Utrecht University

Analyzing and improving the quality of Voronoi Hausdorff morphing using component reduction

MASTER THESIS

Author:
Arjen Simons

Thesis number:
ICA-1393170

First Examiner:
Marc van Kreveld

Second Examiner:
Tim Ophelders

July 7, 2023

Abstract

This thesis introduces a new quantitative metric, used to determine the quality of a morph. This metric is based on the angular change of all enclosed boundary loops and is experimentally shown to effectively distinguish between morphs that are visually appealing and morphs that are worse in terms of visual quality. This thesis also improves the Voronoi and mixed morphs [15] using two approaches. The first approach experimentally shows that a dynamic variable for the mixed morph can result in morphs that better balance all quantitative metrics and results in visually more appealing morphs. The second approach is the introduction of a new abstract morph and its mixed variant, that adjust the Voronoi morph to reduce the number of components present in intermediate shapes. We prove some basic properties on the creation of components in the Voronoi morph and that the new morph also adheres to the Hausdorff property [16]. In an experimental analysis of the new morphs, we record data on the area, perimeter and total angular change development throughout the morph, and the number of holes and components. We show that one of the new morphs performs best on our quantitative analysis and also visually appears most attractive.

1 Introduction

Shape morphing, also known as shape interpolation or shape blending, is the process of gradually transforming a source shape to a target shape over time. There are many different applications for morphing; each with their own techniques. Some are used in the film and game industry to create animations or visual effects. Other techniques are used in medical imaging to generate 3D reconstructions from 2D slices made by MRI- or CT-scanners. In this thesis, morphs are regarded as a function from the interval $[0, 1]$. The parameter operating on this interval can be regarded to as time. The function operates on two input shapes and outputs one shape, such that the output shape at 0 is one input shape and the output shape at 1 is the other input shape. Any shape of the morph at a given time value is referred to as an *intermediate shape*.

Each application has different requirements that determine the quality of a morph. When animating between two poses of an animal, it would make sense that the hind legs in the source shape transform into the hind legs in the target shape. To accomplish this, a lot of techniques use predefined anchor points that define a correspondence between features of the input shapes. When creating a 3D construction from hundreds of 2D slices it would not be practical to rely on user input between each neighbouring slice. It is, however, important that the generated intermediate shapes are anatomically plausible. Generally speaking, all morphs should generate a smooth transition between the source and target shape in such a way that the intermediate shapes preserve the appearance of the input shapes.

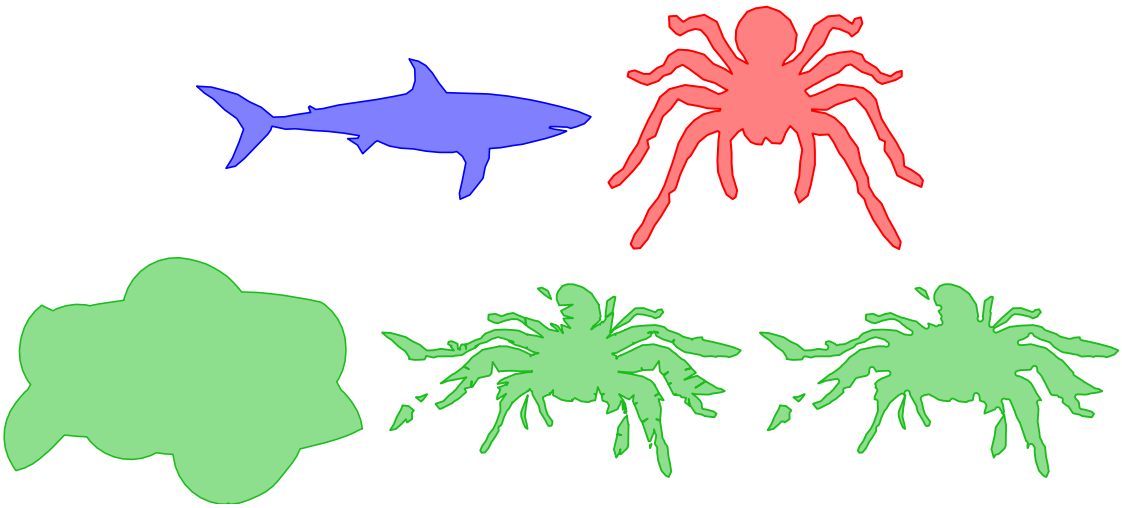

Figure 1: The intermediate halfway morphs of three different morphing methods when morphing between the input shapes at the top. The bottom left shows the dilation morph [16], the bottom middle shows the Voronoi morph [15] and the bottom right shows the mixed morph [15].

This thesis concentrates on a branch in the field of morphing, called abstract morphing, that does not concern itself with the defining features of the input or output shape. A new type of abstract morphing based on the Hausdorff distance is introduced in a recent paper by van Kreveld et al. [16]. The Hausdorff distance is a bottleneck metric that measures how distant two sets are from one another. It is the maximum distance of a set to the nearest point in the other set. Morphs based on the Hausdorff distance are called *Hausdorff morphs* and have the following property: Given two input shapes A and B with a Hausdorff distance of 1 and a time value $\alpha \in [0, 1]$, an intermediate shape can be constructed that has a Hausdorff distance of α to A and $1 - \alpha$ to B. The morph introduced by van Kreveld et al. [16], the *dilation morph* (S_α), is based on the Minkowski

sum with a disk and produces the maximal shape that supports this Hausdorff property. Figure 1 shows how the halfway morph (the intermediate shape at $\alpha = 1/2$) does not have any resemblance with the input shapes as a result of its increased surface area.

In a follow-up paper, de Kogel et al. [15] introduced the *Voronoi morph* (T_α); a Hausdorff morph that resolves the issue of having intermediate shapes with a disproportionate surface area. This morph moves points in the source shape to the closest point on the target shape by a fraction of α and moves points in the target shape to the closest point on the source shape by a fraction of $1 - \alpha$, after which it takes the union of those two sets to get the resulting intermediate shape. Figure 1 shows the Voronoi morph is significantly better in terms of visual quality compared to the Dilation morph. It does however still neglect two main problems: intermediate shapes contain several parts that are fully disconnected from the rest of the shape and the boundary of intermediate shapes contain many slits, which greatly increases the perimeter of the shape. The same paper addressed the latter problem with the introduction of the mixed morph ($M_{\alpha,\varphi}$). It is called the mixed morph because it combines S_α and T_α. It functions the same as T_α, with the addition of an extra step. It first *dilates* T_α using the Minkowski sum, after which it *erodes* the shape using the Minkowski difference with a disc of similar radius. This process is called *closing*, results in the closing of small gaps and holes whilst keeping the rest of the shape intact. This technique does not manage to remove all the slits and still creates extra components at the boundary of Voronoi cells.

This thesis has two main goals. The first is to investigate new ways to more accurately measure the quality of an abstract morph in a quantitative manner. We do this by experimentally investigating different versions of the angular change of intermediate shapes their boundary loops, to determine whether this metric yields distinct outcomes for different morph types and qualities, and to see if this metric can be used to accurately predict which of two morphs will produce better visual results.

The second main goal is to improve the quality of the Voronoi and mixed morph. This is achieved using two different approaches. The first aims to improve the mixed morph, by experimentally testing different parametric inputs for the dilation parameter. We will also investigate the use of a non-fixed dilation parameter that changes over during the morph. In the second approach we present a new Hausdorff morph, based on the Voronoi morph and its mixed variant. This morph is called the Voronoi Component Reduction (VCR) morph. As the name suggest, it aims to reduce the number of components present in intermediate shapes. We experimentally show that this morph produces visually convincing intermediate shapes that preserve more visual features that are present in the input shapes, than the Voronoi morph.

2 Related works

There are many applications for morphing. A big subdomain of morphing is image morphing. This is the process of morphing between two images. The techniques used in image morphing are quite distinct from the techniques used to morph between 2D or 3D shapes, and therefore outside the scope of this thesis. A survey on recent image morphing techniques is presented by Rhora and Kulkarni [17].

A more niche type of morphing is font morphing. Font morphing is used by designers to create new fonts from already existing fonts. It differs from other morphing domains, since it sees morphing as a multidimensional problem, where various aspects of the font, such as the boldness, width and size, can be morphed independently from one another [14].

The focus of this thesis is on the sub-domain that morphs 2D shapes. Most morphing techniques used on 2-dimensional shapes work directly or with some small modifications on 3-dimensional shapes as well. The following subsections discuss different kinds of shape morphing techniques.

2.1 Triangulation

Morphs based on the triangulation of two shapes aim to optimize two problems. (1) Finding a compatible triangulation, and (2) interpolating the positions of the vertices properly.

In order to have a triangulation of two shapes that is compatible, each vertex in one shape must have a corresponding vertex in the other shape. In addition, each edge in one shape connected by two vertices, must have an edge in the other shape that is connected to the corresponding vertices [3]. If this triangulation is generated, each vertex in the source shape can be 'moved' to the position of the corresponding vertex in the target shape using some interpolation method.

Aronov et al. show that a compatible triangulation between two simple polygons with n vertices is always possible [3]. In some cases, this requires the addition of $O(n^2)$ Steiner points inside each polygon. They present two algorithms. The first algorithm creates a triangulation of both polygons P_1 and P_2. These triangulations are then mapped to a convex n-gon, after which the triangulated n-gons are overlayed. At each edge crossing, a Steiner point is added, creating a triangulation that is compatible with both P_1 and P_2. The second algorithm creates a universal triangulation for any n-gon that is independent of the initial polygon. This triangulation has the shape of a spiderweb. It consists of concentric layers of n-gons, in which all the corresponding vertices on a given layer are connected by radial edges. The vertices on the innermost layer are also connected to one central vertex.

The resulting triangulations generated by these two algorithms often have a lot of very elongated triangles. This generally decreases the quality of a morph, especially for textured shapes or images. There are several papers that address this issue, using improved versions of these algorithms.

Alexa et al. [2] refines the first algorithm by improving the initial triangulations of P_1 and P_2. The initial polygons are triangulated using Delaunay Triangulation, which maximizes the minimum interior angle, avoiding elongated triangles. After merging the two triangulations, the resulting triangulation might still contain elongated triangles. To solve this, they move interior vertices and flip interior edges to maximize the minimum angle of the resulting triangulation.

Gotsman and Surazhsky [13] modify the universal triangulation to create a triangulation that allows morphing between two simple polygons of which the intermediate polygons are guaranteed to be simple as well. It also contains fewer Steiner points than the original universal triangulation and does not create any triangles that are very small or skinny. It does so by eliminating redundant Steiner points in the concentric layers of the spiderweb-like structure.

The simplest way to handle the interpolation of vertex position is to interpolate their positions linearly. Linear transformations can lead to a lot of local distortion on textured shapes. It also has difficulties dealing with rotations and typically results in intermediate shapes that are smaller than the source and target shape. Alexa et al. [2] introduce a more complex interpolation model that minimizes local distortion. This method does however require user-defined anchor points.

2.2 Reconstructing 3D shapes

The reconstruction of 3D shapes from 2D slices can also be classified as morphing [1, 4–7]. The information between each slice must be retrieved by interpolating between these slices. This is a relevant problem in the medical world, where MRI scanners are used to create 2D images (slices) of specific parts of the body. A 3D model of the body part can be constructed using morphs between each adjacent slice. Doing so, creates a surface that is not smooth. Each pair of slices is part of a bigger framework of slices. This means that the information of the surrounding slices may also be needed to create a more accurate and smooth reconstruction.

Barequet and Vaxman [6] created a technique that does this using a curve network. The curve network is generated by matching the vertices of each slice with its neighboring slices, creating a flow graph. Using the flow graph a curve network of cubic Bézier curves is created, which in turn is used to construct the mesh. This ensures the slices are connected through smooth curves in the generated 3D mesh. They also introduced a technique that can reconstruct a 3D shape from a collection of colored nonparallel planar cross-sections. The algorithm reconstructs the original 3D space, such that the reconstructed shape also has colored regions that match the input slices [7].

Albu et al. [1] proposed a method that morphs between image slices. This approach morphs between the slices in an iterative manner. Each iteration it uses a parallel deformation process based on conditional dilation. This iterative approach manages to produce smooth 3D volumetric reconstructions of 2D slices. This method only works if all adjacent slices partially overlap.

2.3 Implicit functions

The previously described techniques mostly morph between shapes that are represented using geometric data such as vertex locations and edge connections. This usually requires a vertex correspondence between the shapes. A good matching can be difficult or even impossible to achieve when the shapes have distinct topologies. To circumvent this bottleneck, shapes can be represented as implicit functions.

An example of such an implicit function is a signed distance function (SDF). This function describes the shape S, as an orthogonal distance from a given point x to the boundary S_b of the shape. Points located in the shape, $\{x \in S \text{ and } x \notin S_b\}$ result in a positive value, which will be greater if the point is further away from S_b. When the point lies exactly on the boundary $x \in S_b$, the function returns 0. For any point outside of S, $x \notin S$ the function has negative values [11]. Cohen-Or et al. [12] proposed a method that uses SDFs to morph 3D shapes. It does not simply use linear interpolation between the functions of the source and target shape since this can lead to distorted and relatively small intermediate shapes. To solve this, they use user defined anchor points in the source shape that are linked to anchor points in the target shape. Parts of the shape that are close to an anchor point in the source shape will be moved towards the general area located around the linked anchor point in the target shape.

Turk and O'Brien described another technique based on a variational implicit function. This technique does not create two distinct implicit functions for the source and the target shape. Instead of this, for shapes with n dimensions, they create one implicit function of $n + 1$ dimensions. For 2D shapes, this means that two planes, each for one of the shapes, contain the data constraints of these shapes. These planes are placed parallel to one another in a 3-dimensional space. The implicit function defining a 3D shape that is an interpolation of the 2D input shapes is created using a variational interpolation technique. The resulting intermediate shapes are the contours of slices located between these parallel planes.

2.4 Optimal transportation

The Wasserstein distance (also known as the Earth Mover's Distance) is a distance function that describes the minimal amount of work needed to move all mass of one distribution onto the other distribution. This can be used to morph between shapes in a way that optimally transports the shape its mass. Solomon et al. [18] describe a techniques that can compute the optimal transportation plan in a scalable manner, and is capable of morphing between shapes that have a lot of topological differences.

2.5 Abstract morphing

The Hausdorff distance is not the only shape simularity measure that can be used for abstract morphing. The *Fréchet* distance is another shape similarity measure used to define a morph. It is not clear how this measure can be used to morph between shapes that have different numbers of components and holes, however it is used to morph between shapes with more topological similarities. A *Fréchet matching* proposed by Buchin et al. [10] matches all pairs of points on two curves. This means that an interpolation between these points results in a smooth morph. A problem with this approach is that the outline can self-intersect during the morph. In a later study, Buchin et al. [9] adresses this problem.

3 Preliminaries

Given two sets of points in $\mathbb{R}^2$, A and B, the directed Hausdorff distance from A to B is defined as

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

where d denotes the Euclidean distance, meaning it is the largest distance from all points in A to their closest point on B. The undirected Hausdorff distance is defined as

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

The *dilation morph* defined by van Kreveld et al. [16], where for any time parameter $\alpha \in [0, 1]$

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

where $\oplus$ denotes the Minkowski sum defined as $\{a+b \mid a \in A, b \in B\}$, and D_α is a disc of radius α. S_α has the following properties with respect to the Hausdorff distance: $d_H(A, S_\alpha) = \alpha$ and $d_H(B, S_\alpha) = 1 - \alpha$, assuming $d_H(A, B) = 1$. This means that the Hausdorff distance changes linearly with α. This is called a *Hausdorff property*. Any morph that satisfies this property is known as a *Hausdorff morph*.

The *Voronoi morph* (T_α) defined by De Kogel et al. [15] is defined as

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

where $c(a, B)$ denotes the point on B that is closest to a. In other words, every point in A is moved to its closest point in B by a fraction of α of that distance, and every point in B is moved to its closest point in A by a fraction of $1 - \alpha$. If a point is equidistant to multiple points in the other shape, all options are included. The union of these shapes results in $T(A, B)$. They prove this morph is also a Hausdorff morph.

In the same paper, the *mixed morph* $M_{\alpha,\varphi}$ is defined as

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

where $\ominus$ denotes the Minkowski difference, defined as $A \ominus B := (A^c \oplus B)^c$, where A^c is the complement of A. It first *dilates* T_α by taking the Minkowski sum with a small disc. After that, it *erodes* the shape by taking the Minkowski difference with a disc of the same radius. Performing a dilation proceeded by an erosion with a disc of the same radius is known as *closing*, and can be used to close small gaps and holes while keeping the rest of the shape intact. It is used to remove small slits that are present in the intermediate shapes of the Voronoi morph. This is shown in Figure 1. To make sure it is a Hausdorff morph, the intersection with S_α is taken.

4 Experimental setup

All three major sections will be tested on the same data sets and similar implementations. The implementation details and data sets will be described in the following subsections.

4.1 Implementation details

The morphs are implemented in C++, using Boost[1] to calculate intersections and unions of polygons and lines, Voronoi diagrams and Minkowski sums. A Voronoi diagram of a polygon may contain curves. Our implementation approximates these curves with line segments, which can result in a very small error. This error should not manifest as a noticeable difference in the resulting morphs. Another error occurs from the final step in the implementation, where all slices partitioned by the Voronoi diagrams are combined using a tiny dilation to ensure neighbouring slices are combined properly. This can impact the perimeter and number of component in intermediate shapes when two slices merge while there should be a tiny distance between them. This effect will be more noticeable as α approaches 0 and 1. The area will only slightly be affected by this.

4.2 Data sets

The tests are performed on three data sets. The first data set contains a collection of 9 outlines of animals taken from [8]. Each animal polygon only consist of one component. The second data set is a collection of 13 European country outlines from the Thematic Mapping World Borders data set:[2] Austria, Belgium, Croatia, Czechia, France, Germany, Greece, Ireland, Italy, the Netherlands, Poland, Spain and Sweden. Except for Austria, Belgium Czechia and Poland, all country polygons are comprised of multiple components. The last data set contains a small collection of letters, traced as polygons, taken from [15]. Three pairs of words will be used (wish/luck, kick/stuff, try/it), and the letters f, i and u in a serif and a sans serif font.

4.3 Experimental setup

None of the morphs are translation- or scale-invariant. Therefore the resulting intermediate shapes differ depending on the initial position and scale of the input shapes. In all experiments, input shapes from the animal and country data sets are scaled to have the same surface area and are translated to have a common centroid. Input shapes from the letter data set, will not be scaled to have the same surface area; font size is used instead.

The assumption is made that an ideal morph linearly interpolates the area and perimeter between those of the input shapes. For each experiment, the ratio between the measured area and perimeter and these 'ideal' values is given. The number of components and holes are discrete and are directly recorded. Since the total angular change—discussed in the next section—does not have a clear ideal value, these values are also directly recorded.

The type of experiments differ per section and are therefore discussed separately in each section.

[1] https://www.boost.org

[2] https://www.thematicmapping.org/downloads/world_borders.php

5 A new quality metric

An ideal morph is a morph that visually represents the input shapes most accurately based on the time variable α. When $\alpha = 1/2$, the intermediate shape should partially represent both input shapes equally without showing details that do not visually represent any parts of the input shapes. When $\alpha = 1/4$, the intermediate shape should mostly represent the source shape, while still resembling some of the target shape.

De Kogel et al. [15] use both quantitative and qualitative metrics to determine the quality of a given abstract morph. The quantitative metrics include the change in surface area, perimeter, number of components and holes, for different values of α.

An ideal morph does not create extra components that do not visually resemble any details of the source or target shape. In most cases, any component created during the morph that continues to merge into another component, creates extra detail that does not contribute to any visual resemblance of the input shapes. Morphs between some input shapes might require extra components to be created in order to not distort the input shapes too much. If we morph between the shape of the letter C and a mirrored C, the ideal morph might split the letter in half horizontally. However, for most inputs, the ideal morph should probably not have more components than either one of the input shapes during any stage of the morph.

The number of holes in an ideal morph might exceed that of either one of the input shapes. If the hole from the source shape is close to that of the target shape, an ideal morph, might translate and scale one hole to the other. In that case, the number of holes should not increase. However, if the input shapes both have a hole on opposite sides, an ideal morph might open and close the holes at the same time during the morph. This causes the number of holes to exceed that of the target and source shape. Therefore, the number of holes in an ideal morph should not exceed the number of holes in the source and target shape combined.

Ideally, the surface area and perimeter of intermediate shapes in a morph scale linearly with α between that of the source and target shape. If the surface area of intermediate shapes exceeds that of a linear interpolation, it is likely that the shape loses resemblance with the source and target shape. An example of this is the dilation morph shown in Figure 1. The perimeter of intermediate shapes shows us whether details of the input shapes might be lost, if it is smaller than a linearly interpolated perimeter or whether extra detail is created that does not resemble any of the input shapes if the perimeter is larger. An example of the latter case is the Voronoi morph shown in Figure 1.

In this section, a new qualitative metric will be explored to get more insight in the quality of a given morph. This metric is based on the angular change of the enclosed boundary loops. Investigating the angular change of intermediate shapes can inform us about the conservation of concave details of the input shapes during a morph.

The following subsections explore different implementations of the angular change metric. For each implementation, the dilation, Voronoi and mixed morphs for two pairs in the animal dataset is computed and examined: shark→spider and cat→dog. In all experiments, the angular change implementation is measured for α values starting at 0, in increasing steps of $\frac{1}{100}$. The parameter φ of the mixed morph, is set to 0.02 for all experiments in this section as a starting point, because this value for φ is also used in the paper that introduced the mixed morph [15].

5.1 Total angular change

The most basic angular change metric that will be explored is the sum of the angular change of each closed boundary loop. If the shape is composed of n convex components and holes, then the

total angular change will be $n2\pi$. The lowest value for the total angular change of a given shape is 2π, because concavities will always increase the angular change of a closed boundary loop.

When morphing between a source and target shape, the intermediate shapes should gradually lose resemblance with the source shape and gain resemblance with the target shape as α increases. At $\alpha = 1/2$, the angular change should reflect that concavities from both input shapes are preserved. Therefore, the expected angular change at $\alpha = 1/2$ is larger than that of the source and target shape. The angular change in an ideal morph should gradually increase as α grows and peak around $\alpha = 1/2$, after which it should gradually decrease to that of the target shape. It should not exceed the angular change of the source and target shape combined, since that would mean concavities or components not present in the input shapes appear during the morph.

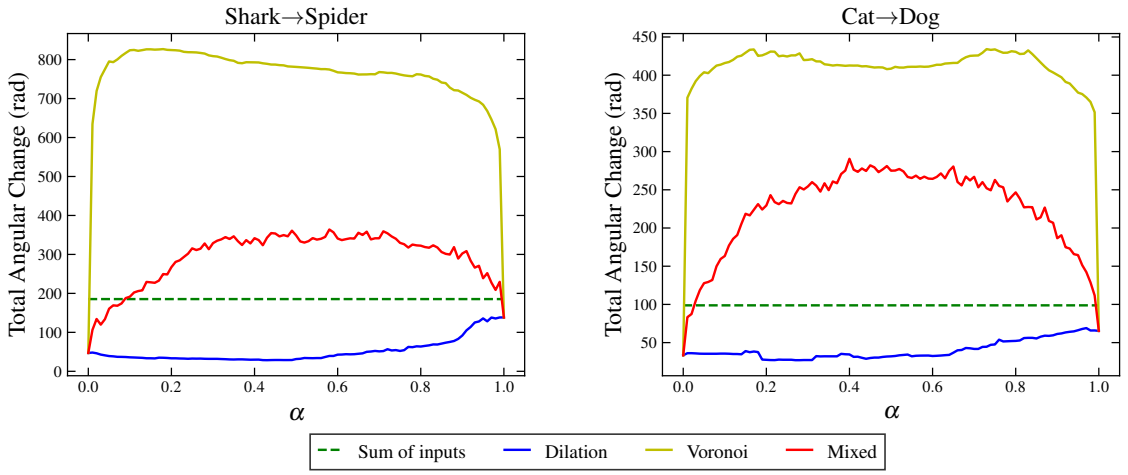

Figure 2: The total angular change in radians as a function of α when morphing from the shark to the spider on the left, and from the cat to the dog on the right.

The results shown in Figure 2 are in line with the predictions made at the beginning of this section. The dilation halfway morph between the shark and spider is less concave than the shark and especially less concave than the spider (see Figure 1). This is reflected in the total angular change, which turns out to be lower than either of the input shapes. The cat→dog dilation morph also has a decrease in the total angular change on the Dilation morph.

The total angular change of the Voronoi morph instantly increases a lot as α changes from 0 to $\frac{1}{100}$, after which it seems to only slightly change until it decreases again when α changes from $\frac{99}{100}$ to 1. This rapid change comes from the slits and components that arise when points on opposite sides of a Voronoi edge move in different directions. Causing them to split or merge at $\alpha = 0$ and $\alpha = 1$. The extra concave details created in the Voronoi morph are clearly represented in the total angular change, as it is more than four times the sum of the inputs' total angular change for most of the intermediate shapes.

The mixed morph *closes* small slits and merges components when their distance falls below 2φ. This is reflected in the total angular change; it gradually increases and peaks around $\alpha = 1/2$. This is when slits are at their peak size and the distance between components is largest. Large slits that do not represent features in the input shapes are still present in the halfway morph. This is reflected in the total angular change, which is quite a bit lower than that of the Voronoi morph, but still almost double the sum of the inputs' total angular change in the shark→spider morph and almost three times the sum of the inputs in the cat→dog morph. The curvature of the total angular change over α for the mixed morph is also less smooth than that of the Voronoi and dilation

morphs. This mostly comes from components and slits that suddenly appear or disappear during the morph.

5.2 Average angular change

The total angular change is quite heavily influenced by the number of components that are created during the morph, because every components adds at least 2π to the total angular change of a given shape. A way to reduce this effect is to take the average angular change of all the enclosed boundary loops to determine the quality of a morph.

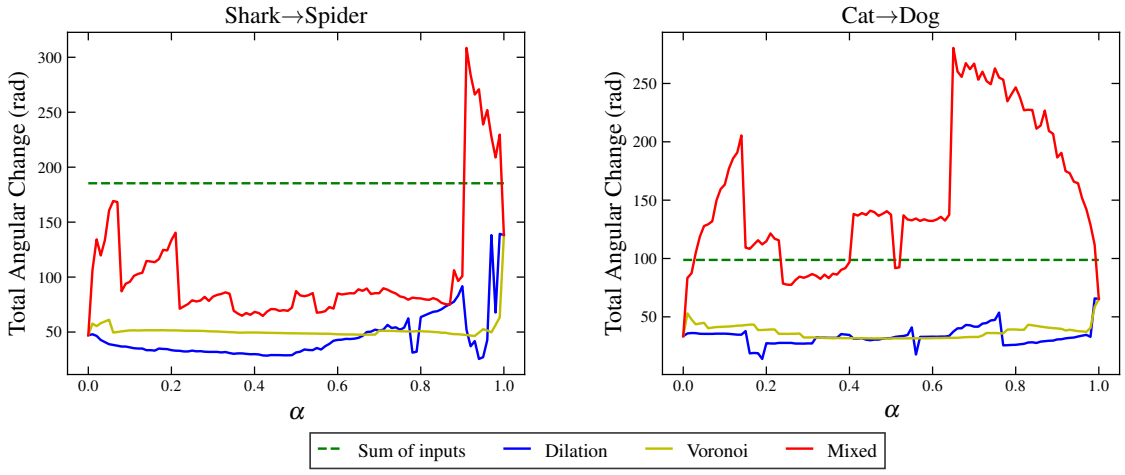

Figure 3: The average angular change in radians as a function of α when morphing from the shark to the spider on the left, and from the cat to the dog on the right.

The results shown in Figure 3, indicate that average angular change of the shark→spider dilation morph is the same as the total angular change of the shark→spider dilation morph shown in Figure 2, up until $\alpha = 0.78$. After that, the average angular change has some rapid changes. After $\alpha = 0.78$, holes start to appear. These holes do not affect the total angular change too much because they are quite convex. Therefore, when the average is taken, the angular change suddenly drops when a hole appears.

The same effect can be observed on the average angular change in the mixed morph. It changes significantly between small intervals of α when components merge or disconnect. During the mixed morph, several holes and components disappear or appear as they fall below or outside of the 2φ range (the distance it takes for two disconnected parts to be merged).

Contrary to the dilation and mixed morph, the average angular change of the Voronoi morph as a function over α is quite smooth. The number of components in the Voronoi morph only change at $\alpha = 0$ or $\alpha = 1$ [15]. Therefore, taking the average of all components between small intervals of α, does not suddenly change the average angular change very much. The average angular change is however quite low, compared to the mixed morph and seems to be quite similar to that of the dilation morph. This is a result of the number of components present in the Voronoi morph. Most small components are quite convex. This means they do not attribute that much to the total angular change. When taking the average, this is not taken into account. Therefore, these convex components strongly lower the average.

5.3 Perimeter weighted angular change

The total and average angular change do not take into account the fact that some components contribute more to the visual quality of a morph than other components. When observing an intermediate shape that is comprised of two components – one that is very large and one that is very small – the details on the large component contribute most to the visual quality of the intermediate shape. The perimeter weighted angular change takes this into account by weighting each closed boundary loop with its perimeter divided by the total perimeter of the intermediate shape. This way, small components contribute less to the computed angular change.

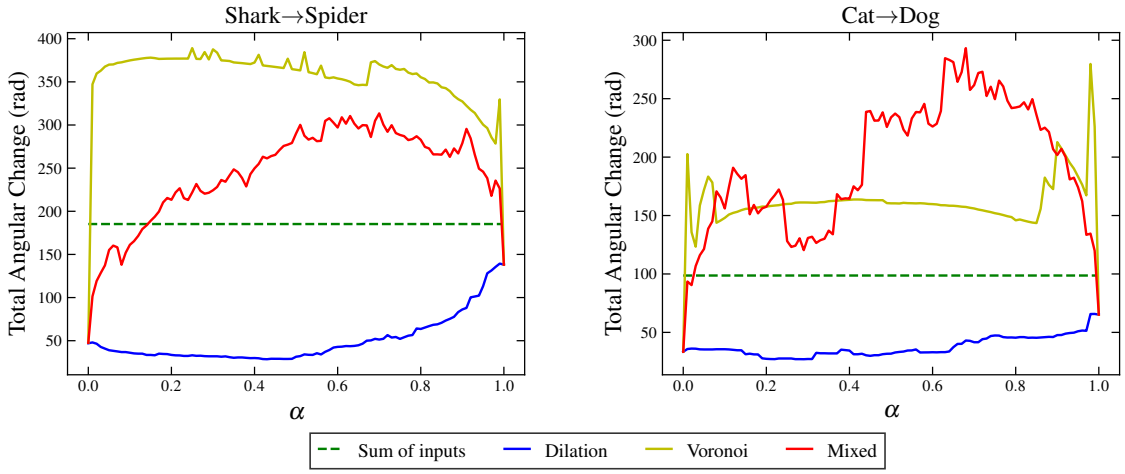

Figure 4: The perimeter weighted angular change in radians as a function of α when morphing from the shark to the spider on the left, and from the cat to the dog on the right.

he results shown in Figure 4 are quite similar to that of the total angular change. The perimeter weighted angular change of the shark→spider Voronoi morph seems to be the almost identical in gradient to that of the total angular change; it is only scaled down on the y-axis. Both tests – most noticeably in the cat→dog Voronoi morph – seem to have some abrupt changes in the angular change. This seems to be a result of the tiny dilation that is used to merge all slices that are partitioned by the Voronoi diagrams of the input shapes. This causes some components to connect when they are very close while there should be a tiny distance between them. Because slits are thinnest close to $\alpha = 0$ and $\alpha = 1$, this effect is most noticeably close to those α values.

The mixed morph's perimeter weighted angular change differs more from the total angular change than the dilation and Voronoi morphs their perimeter weighted angular changes. It decreases more rapidly as it approaches $\alpha = 1$. The impact of components and holes appearing/disappearing is quite noticeable in the cat→dog mixed morph, by the rapid increase/decrease of the angular change at small intervals of α.

5.4 Comparison

The average angular change does not indicate the quality of a morph really well. It mostly shows whether components or holes appear/disappear during the morph. The average angular change of the Voronoi and dilation morphs are quite similar despite being visually very distinct.

The perimeter weighted angular change, indicates the quality of a morph in a clearer manner. The shark→spider test (Figure 4) clearly shows that the Voronoi morph's intermediate shapes have more concavities than those of the dilation and mixed morphs. It is however still affected quite

heavily when components or holes appear/disappear. This is most noticeable in the cat→dog test shown in Figure 4, where the results indicate that most intermediate shapes at $\alpha > 0.4$ in the mixed morph are more convex than those in the Voronoi morph. Visually this is however not intuitive, as can be seen in Figure 1.

The total angular change appears to most accurately indicate the quality of a given morph. Each morph shows distinct values that characterise visual attributes of the morphs (Figure 2). The low angular change of the dilation morph is in line with its convex nature and the high values from the Voronoi morph are in line with all the concave slits that are present in intermediate shapes. The mixed morph is visually the most balanced in terms of convexity. The total angular change metric captures this, as the values are in between that of the Voronoi and dilation morphs. Therefore, the total angular change is used in all experiments in the next sections.

6 Mixed morph improvements

The mixed morph has an input-parameter φ that denotes the radius of the disc used for the closing operator. De Kogel et al. [15] experiments with the mixed morph using one value for this parameter: $\varphi = 0.02$. Higher values of φ can result in less components, but it also increases the surface area resulting in more details of the input shape being lost during the morph. This section explores different values for φ and explores φ as a function where the output of φ changes with α.

6.1 Different φ values experiments

In these experiments the mixed morph with different parametric inputs for all pairs of animals and all pairs of countries in the animals and countries data sets are computed. The word and letter pairs from the letter data set are also computed. In each experiment, the surface area, perimeter, total angular change, number of components and number of holes of the morph will be measured for α values starting at zero and increasing in steps of 1/8. Several values of φ will be tested: 0.01, 0.02, 0.03, 0.05, 0.1 and 0.2.

6.1.1 Results

Tables 1 and 2 show a summary of the surface area and perimeter measurements. Table 3 contains a summary of the number of components and holes for the animals data set; the other data sets are excluded because the inputs have different numbers of components.

Figure 5 shows that the dilation parameter has a big impact on the surface area of intermediate shapes. On both the animals and countries data sets, $\varphi = 0.02$ and $\varphi = 0.03$ best fit a linear interpolation of the surface area between the input shapes.

The perimeter, as shown in Figure 6, is also greatly affected by the dilation parameter. $\varphi = 0.05$ stays closest to 1.0 for any value of α. On the animals data set the perimeter grows a bit larger than 1.0, while it stays lower than 1.0 in the countries data set. In both data sets the perimeter is smaller than 1.0 close to $\alpha = 0$ and $\alpha = 1$ when $\varphi = 0.05$. This is a result of slits growing as α approaches 0.5. If α is close to 0 or 1, the slits are thin, and will therefore be *closed*. $\varphi = 0.03$ also does not diverge too much from the *ideal* linear interpolation.

The number of components differs most between small values of φ. Table 3 shows that after $\varphi = 0.1$, there are not as many components left to be merged with other components when increasing the dilation parameter. For the tested values of φ, the number of components is smaller when φ is larger. This is not the case for the number of holes. The number of holes is largest when $\varphi = 0.05$. When φ increases, it can *close* small holes, however, larger values of φ can also create

holes, when merging the ends of two bulges. An example of this can be seen in Figure 8, where only the morph with $\varphi = 0.05$ has a hole in it.

Figure 7 shows that there is a clear difference in the total angular change between different values of φ. As α increases, the total angular change decreases for the tested values of φ. $\varphi = 0.1$ has the best fit in terms of total angular change on the animals data set since it peaks close to the sum of the inputs. On the countries data set, $\varphi = 0.03$ peaks very close to the sum of the inputs, but on the animals data set it is significantly larger. $\varphi = 0.05$ has a more balanced total angular change between the two data sets. It does not however seem to have a good fit for either of the data sets. On the Animals data set it exceeds the sum significantly, and on the countries data set it peeks at 2/3 of the sum of the inputs.

Overall, $\varphi = 0.03$ seems to result in the best morphs. In terms of surface area, perimeter is performs close to the ideal linear interpolation. In terms of total angular change it seems perfect on the countries data set, but peeks too high on the animals data set. In terms of components $\varphi = 0.05$ is better. This value of φ also has reasonable results in surface area and perimeter metrics. It does however also create more holes that have a significant impact on the visual quality of the morph. This can be seen in Figure 8, where the morph with $\varphi = 0.05$ at $\alpha = 0.75$ has three holes that are quite large. In this Figure the merging of components by the closing operator can also be observed. Merging components that are not very close using the closing operator causes a bulge that does not look very pleasing. The morph at $\varphi = 0.05$ at $\alpha = 0.75$ merges a component with the rump of the spider. This creates a bulge that does visually does not represent any details in the source or target shape.

	Animals		Countries		Text	
φ	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
0.01	0.982	0.020	0.978	0.035	1.029	0.030
0.02	0.991	0.019	0.994	0.039	1.033	0.030
0.03	1.002	0.019	1.011	0.049	1.039	0.034
0.05	1.024	0.027	1.046	0.081	1.052	0.044
0.1	1.078	0.064	1.121	0.171	1.097	0.084
0.2	1.169	0.136	1.215	0.284	1.239	0.290

Table 1: The distributions of normalized surface areas for each tested φ value for all tested values of α, in the mixed morph, separated by experiment category

	Animals		Countries		Text	
φ	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
0.01	1.277	0.194	1.249	0.239	1.170	0.213
0.02	1.175	0.151	1.122	0.180	1.141	0.186
0.03	1.107	0.126	1.045	0.152	1.117	0.168
0.05	1.019	0.099	0.949	0.137	1.083	0.139
0.1	0.898	0.098	0.837	0.155	1.029	0.080
0.2	0.800	0.139	0.764	0.185	0.950	0.067

Table 2: The distributions of normalized perimeters for each tested φ value for all tested values of α, in the mixed morph, separated by experiment category

	Components		Holes	
φ	Mean	Std. Dev.	Mean	Std. Dev.
0.01	6.702	3.492	0.163	0.439
0.02	4.976	2.955	0.214	0.514
0.03	3.940	2.499	0.274	0.558
0.05	2.710	1.866	0.349	0.695
0.1	1.671	1.074	0.290	0.624
0.2	1.155	0.373	0.238	0.584

Table 3: The distributions of the number of components and holes for each tested φ value for all tested values of α except 0 and 1. Only the animal data set is included, as these shapes only have one component and no holes.

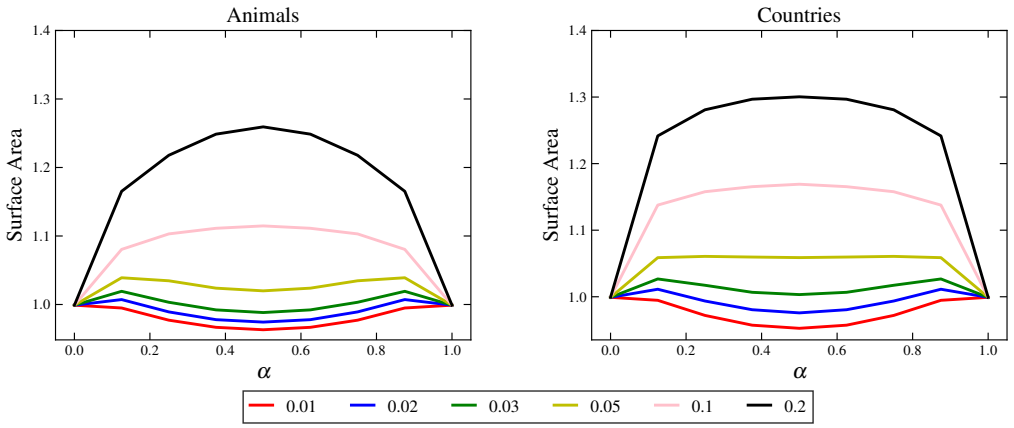

Figure 5: The normalized average surface area over all experiments as a function of α, for both the animals and countries data sets.

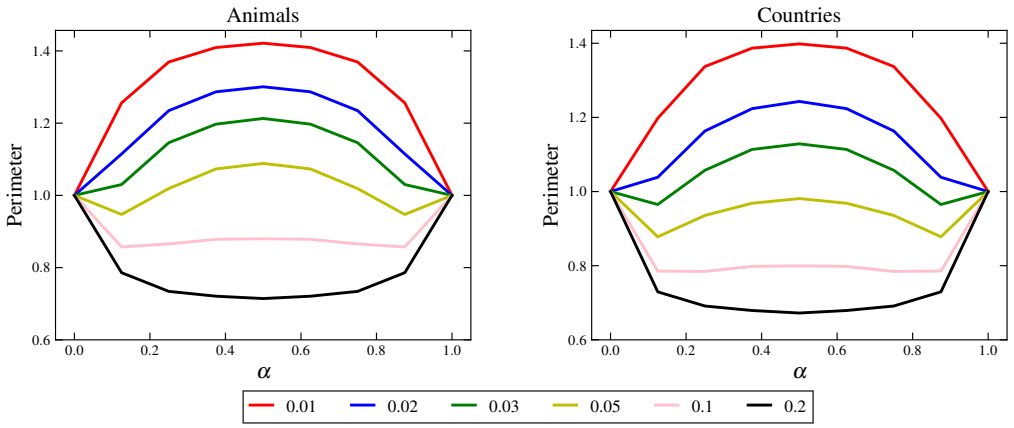

Figure 6: The normalized average perimeter over all experiments as a function of α, for both the animals and countries data sets.

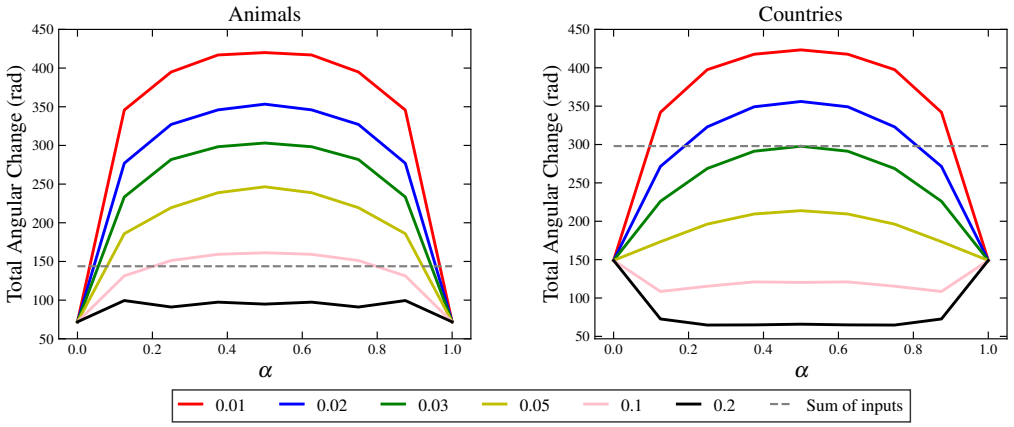

Figure 7: The average total angular change over all experiments as a function of α, for both the animals and countries data sets. The dashed line indicates the average sum of the input shapes their angular change for all experiments.

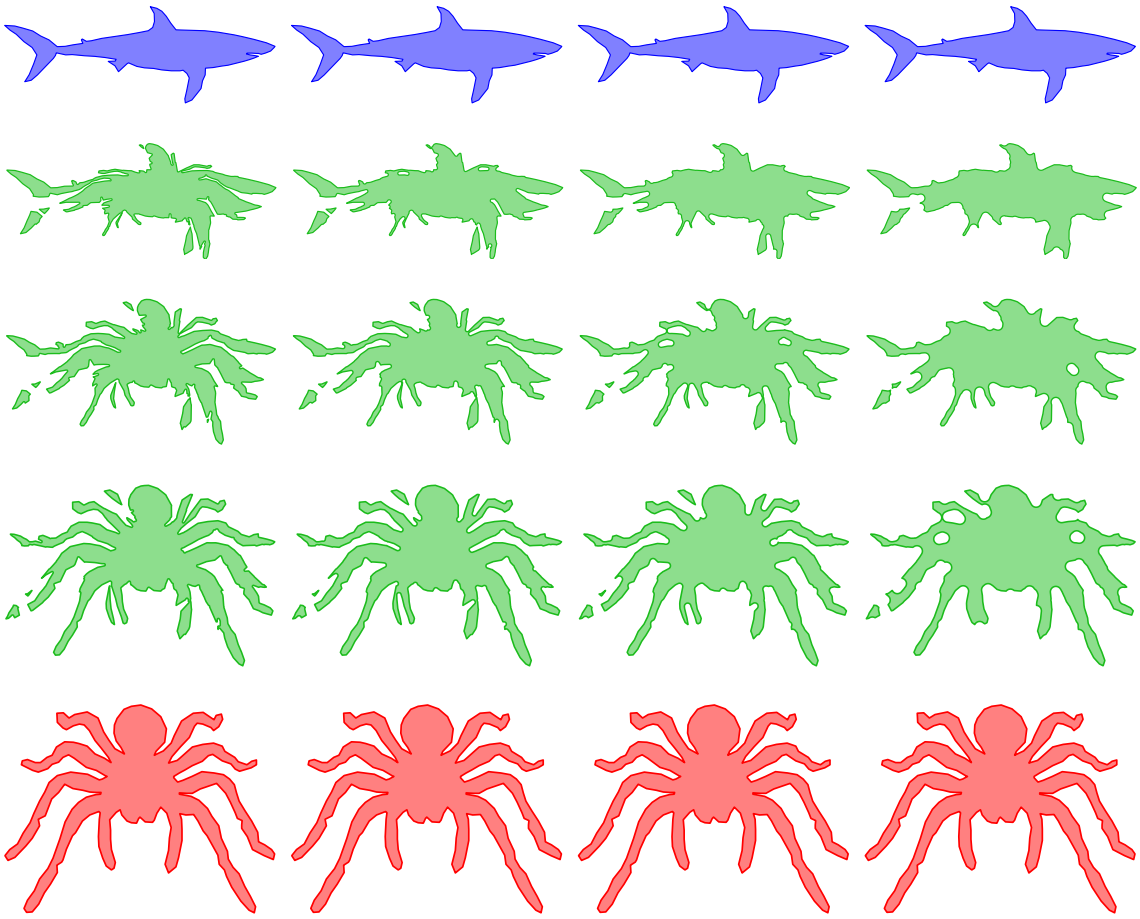

Figure 8: Intermediate shapes for $\alpha \in \{0, 1/4, 1/2, 3/4, 1\}$ when morphing between outlines of a shark and spider. The columns show the mixed morph with different values for φ: 0.01, 0.02, 0.03 and 0.05 from left to right.

6.2 φ as a function experiments

In the previous section, φ was constant for all values of α. Meanwhile, attributes of intermediate shapes such as slits are very thin α values close to 0 and 1. This means that a small value of φ can be enough to *close* these slits. In Figure 6, it can be seen that the high value of φ (0.05) has a significant impact on the perimeter for α values close to 0 and 1.

To overcome this, φ can be a function of α, where φ is smaller near $\alpha = 0$ and $\alpha = 1$ and larger near $\alpha = 1/2$. In this section we will test the following simple piecewise linear function for φ:

$$f(\alpha) = \begin{cases} 1.475\alpha + 0.005 & \text{if } \alpha < 0.2 \\ 0.03 & \text{if } 0.2 \leq \alpha \leq 0.8 \\ -1.475\alpha + 1.48 & \text{if } \alpha > 0.8 \end{cases} \tag{1}$$

At $\alpha = 0$, φ starts at a tiny value (0.005) to make sure no instant slits are created that are present in the Voronoi morph. When $\alpha < 0.2$, φ linearly grows until it reaches 0.03, after which it stays constant up until $\alpha = 0.8$. For $\alpha > 0.8$, φ linearly decreases to 0.005 as α approaches 1.0. The $\alpha = 0.2$ and $\alpha = 0.8$ cutoff points are chosen based on preliminary experimentation.

To test φ as a function over α, two experiments are executed. In the first, φ will be set to 0.03, as this turned out to be the most reasonable value for φ of all tested values. In the second, φ will be set to the function shown in Equation 1. In both experiments the mixed morph for the word and letter combinations and all combinations in the animal and country data sets are computed. In each experiment, the surface area, perimeter, total angular change, number of components and number of holes of the morph will be measured for α values starting at zero and increasing in steps of 1/8.

6.2.1 Results

Figure 9 and 6 show that when φ is not fixed, it can better fit a linear interpolation between the surface area and perimeter of the input shapes. With a fixed φ at 0.03, the perimeter and surface area instantly change significantly at the start and end of the morph. φ represented as a piecewise linear function eliminates this sudden change.

On the angular change, the opposite effect can be observed. Figure 11 shows that the angular change at $\alpha = 0.1$ exceeds that of α values closer to $\alpha = 0.3$ when representing φ as a the function shown in Equation 1. The fixed $\varphi = 0.03$ value results in a *smoother* change of the angular change over α. Small details stay intact when φ is small, while new details can arise as well. Therefore, it is not necessarily a bad development of the angular change.

In terms of components, Table 4 shows the fixed value $\varphi = 0.03$ performs better. At the start and end of the morph some components might not be merged when $\varphi < 0.03$. In terms of holes, φ represented as a function performs better. It is however likely that components present at the start and the end of the morph when φ is represented as a function, will turn out to be components at $\alpha = 1/2$, because these components only exists at small alpha values if the distance to their target in the other shape is large, which means they move *away*, from other components more rapidly.

Representing φ as a function shows promising results, as it performs better in terms of area, perimeter and the number of holes. Meanwhile, it performs worse in terms of the number of components. It is however, likely that these components are also present with a fixed φ value at other stages of the morph. For this reason, φ represented as the function shown in (1), will be used as a benchmark for the mixed morph when comparing it to other morphs in the next section.

	Components		Holes	
φ	Mean	Std. Dev.	Mean	Std. Dev.
function	4.210	2.480	0.337	0.663
0.03	3.940	2.499	0.274	0.558

Table 4: The distributions of the number of components and holes for each tested $\varphi = 0.03$ and φ as the function in Equation (1). Only the animal data set is includes, as these shapes only have one component and no holes.

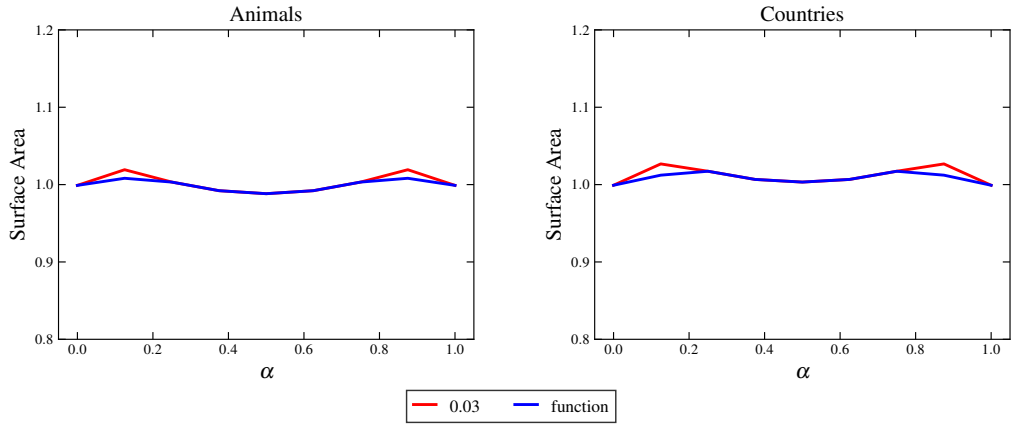

Figure 9: The normalized average surface area over all experiments as a function of α, for both the animals and countries data sets.

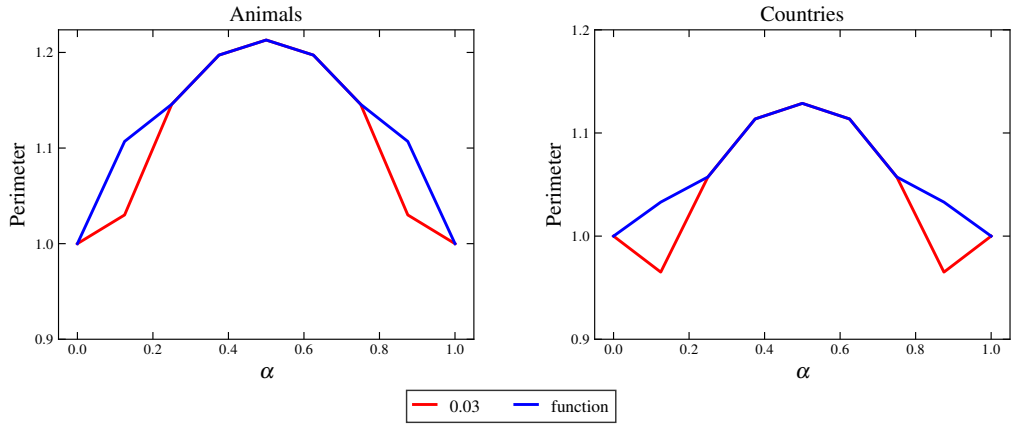

Figure 10: The normalized average surface area over all experiments as a function of α, for both the animals and countries data sets.

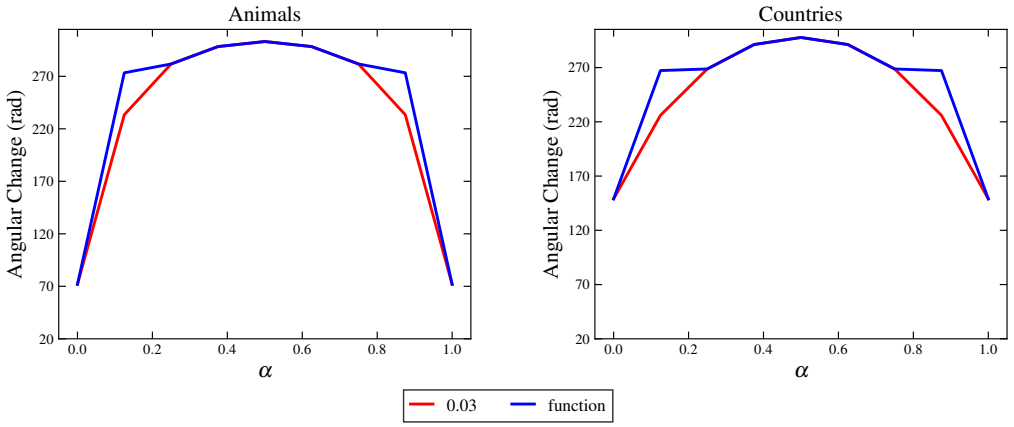

Figure 11: The average total angular change over all experiments as a function of α, for both the animals and countries data sets. The dashed line indicates the average sum of the input shapes their angular change for all experiments.

7 A new morph

The number of extra components in intermediate shapes has a big impact on the quality of the morph. The previous section showed how increasing the dilation parameter in the mixed morph reduces the number of components created. The value of φ needed to remove most components decreases the quality of the morph in terms of the surface area and perimeter quality metrics.

In this section a new Hausdorff morph is introduced. The new morph is based on the Voronoi morph. It tries to lower the number of extra components created in the Voronoi morph. The Voronoi morph linearly moves every point from one input shape A to its closest point in the other input shape B. Both input shapes are partitioned by the Voronoi diagram of the other shape. Each part of the shape located inside a Voronoi cell is moved to the site of that Voronoi cell. This site can be a vertex, edge or face of the polygon that represents the shape [15]. Moving a part to its corresponding site can create an extra component when its neighbouring Voronoi cells move in different directions. The new morph identifies when a component is created and if possible, moves it to sites of a neighboring Voronoi cell.

The first subsection investigates when and how components are created in the Voronoi morph.

7.1 Voronoi morph Components

De Kogel et al. [15] show that the number of components in a Voronoi morph $\#C(T_\alpha)$ is constant for every value of α in the range $(0, 1)$. This means that components only appear and disappear when changing α between 0 or 1 and another value of α.

Let $V(A)$ be the Voronoi diagram of the vertices, open edges and interior components of A and let $\text{Par}(A, B)$ be input shape A partitioned into regions by $V(B)$. From now on we shall refer to these partitioned regions as slices. This means all points in a slice $S \in \text{Par}(A, B)$ have a closest point, edge or interior component in B, which can be described as the site of the Voronoi cell in $V(B)$ in which S lies.

We define the directed Voronoi morph as:

$$\vec{T}_\alpha(A, B) := \{a + \alpha(c(a, B) - a) \mid a \in A\}$$

where $c(a, B)$ denotes the point on B closest to a. It moves any point in A to its closest point in B by a fraction of α. Note that the shape at $\alpha = 1$ does not have to be equal to B. Only if B fully overlaps with A, this will be the case. If they it does not overlap on some parts, A will just morph to the boundary of B.

We can prove that the number of components in the directed Voronoi morph, will not exceed the number of components in A, $\#C(A)$, when B is convex.

Lemma 1. *Let B be a convex shape. Then $\#C(\vec{T}_\alpha(A, B)) = \#C(A)$ for any $0 \leq \alpha \leq 1$.*

Proof. We first show that any slice $S \in \text{Par}(A, B)$ does not lose connectivity during the morph. After that, we show that components consisting of multiple slices, also do not lose connectivity.

Consider the site of S its corresponding Voronoi cell to be an interior component. The initial position of S overlaps with the target shape. Therefore, the slice will not move when changing α, and will be connected for any α.

Now consider the site of S its corresponding Voronoi cell to be a line segment L. The vector $c(p_1, B)$, where $p_1 \in S$ will be perpendicular to L, and parallel to any vector $c(p_2, B)$, where $p_2 \in S$. Because of this and the fact that every point $p \in S$ is moved linearly towards L with respect to α, S will be connected for any α.

Now consider the site of S its corresponding Voronoi cell to be a point P. Assume two points $p_1, p_2 \in S$ that are not collinear with P. Since P is the closest point in B of p_1 and p_2, the vectors $c(p_1, B)$ and $c(p_2, B)$ intersect at P. Therefore p_1 and p_2, will move closer to each other if α increases. Now assume p_1, p_2 and P are collinear. In this case the vectors $c(p_1, B)$ and $c(p_2, B)$ have the same direction. This means that p_1 and p_2 will also move closer to each other if α increases. Since all points in S move closer to each other when α increases, S will be connected for any α.

Now consider two slices $S_1, S_2 \in \text{Par}(A, B)$ that are connected at $\alpha = 0$ and do not overlap with B at $\alpha = 0$. Every point $p \in S \in \text{Par}(A, B)$ moves in a straight line to its closest point in B. This means that S_1 and S_2 must be connected through the edge e that divides their corresponding Voronoi cells. Let $V_{exterior}(B)$, be the part of $V(B) \notin B$. Because B is convex, two Voronoi cells in $V_{exterior}(B)$ can only be adjacent if the site of one cell is a line segment and the other is a point. Therefore if S_1 morphs towards a line segment, S_2 has to morph towards a point, and if S_1 morphs towards a point, S_2 has to morph towards a line segment. If the site of a Voronoi cell $v \in V_{exterior}(B)$ is a line segment, the edges of v will be perpendicular to that line segment. This means that any point on the edge of v moves along with this edge. If the site of a Voronoi cell $v \in V_{exterior}(B)$ is a point P, the edges of v are pointed directly to P. Therefore any point on an edge of v moves along with that edge. Because S_1 and S_2 are connected through an edge, they both have overlapping points on that edge. These points will move in the same direction along the edge with respect to α. Therefore S_1 and S_2 stay connected for any α.

Now consider two slices $S_3, S_4 \in \text{Par}(A, B)$ that are connected at $alpha = 0$ and where only S_3 overlaps with the interior of B at $\alpha = 0$. The part connecting the S_3 and S_4 has to be at the boundary of B. This is either a vertex or an edge. This means that all points in S_4 will move towards the closest point on that edge or vertex. Thus, points already on the edge or vertex, that are connected with S_3 will not move as α changes. Since S_3 overlaps with B at $\alpha = 0$, it will not move as α changes. Therefore S_3 and S_4 stay connected for any α.

In the last case we consider two slices $S_5, S_6 \in \text{Par}(A, B)$ that are connected at $\alpha = 0$ and both overlap with B at $\alpha = 0$. Since both shapes overlap with B, S_5 and S_6 will not move if α changes. Therefore, S_5 and S_6 stay connected for any α.

Since slices do not lose connectivity for any α, and any two slices that are connected at $\alpha = 0$ also do not lose connectivity for any α when B is convex, the connectivity of any component stays

the same for any $0 \leq \alpha \leq 1$. Furthermore, as proven by de Kogel et al. in Lemma 3 of [15], components can never appear. Thus, the statement in this lemma follows. □

Lemma 2. *Let A and B, be two convex polygons, each consisting of one component. Then $\#C(T_\alpha(A, B)) = 1$ for any $0 \leq \alpha \leq 1$.*

Proof. Lemma 1 demonstrates that all points in A will be connected to each other for all values of α. It also demonstrates that all points in B stay connected because $T_\alpha(B, A)$ results in the same morph; only parameterised in reverse.

Let there be a closest pair of points a, b such that $a \in A, b \in B$. This means that $c(a, B) = b$ and $c(b, A) = a$. Therefore, a is translated to b, by a fraction of α in a straight line l, and b is translated to a by a fraction of $1 - \alpha$ on l. Therefore $b = a$ for any $0 \leq \alpha \leq 1$. This means that A and B will be connected at any state of the morph, because two convex polygons always have a closest point. Thus, $\#C(T_\alpha(A, B)) = 1$ for any $0 \leq \alpha \leq 1$. □

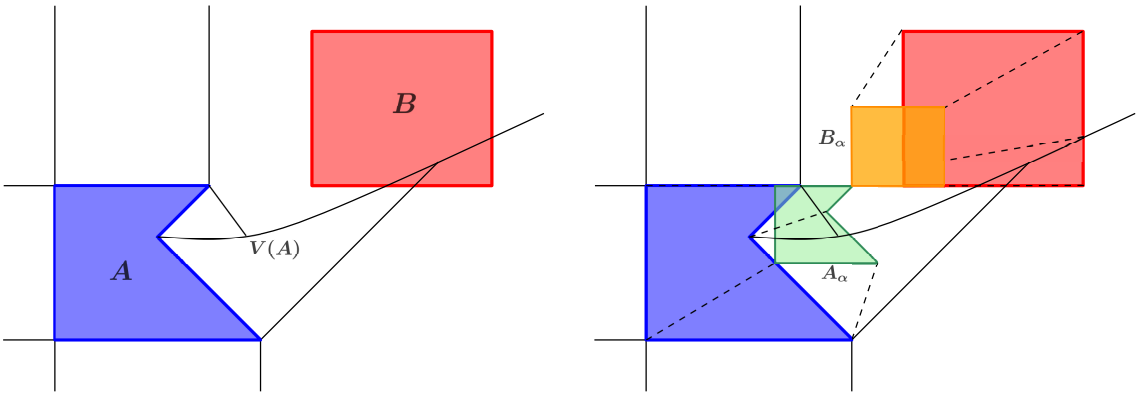

Figure 12: On the left, B is partitioned by the Voronoi diagram $V(A)$ of A. On the right, each partitioned part of B, shown in orange, is scaled towards the closest point in A by a factor $1 - \alpha$, and the morphed A, shown in green, is scaled towards the closest point in B (bottom left vertex of B) by a factor of α.

These proofs show that convex parts of the input shapes are not responsible for any extra components created during the morph. Next we prove that a component is created when it is split by a Voronoi edge originating from a concave vertex if it contains no point that are part of the closest point mapping between A and B. An example of this is shown in Figure 12.

We define E as a set of boundary edges and exterior edges in $V(A)$ that arises from a concave vertex $v \in A$ and partitions B into two sets of slices $S_r, S_l \subset \text{Par}(B, A)$, in such a way that S_r contains all slices on the right of E and S_l contains all slices to the left of E.

Lemma 3. *Let there be no point $p_r \in S_r$ such that $c(a, B) - a = p_r$ and $a \in A$, then S_r will be disconnected for any $0 < \alpha < 1$.*

Proof. For this proof, we first demonstrate that S_r and S_l are not connected during the morph. After that, we demonstrate that S_r is not connected to any part of A during the morph.

Every point $p \in \text{Par}(B, A)$ moves in a straight line toward to its closest point in A. Therefore, two slices $S_1 \in S_l$ and $S_2 \in S_r$, can only be connected at any $0 < \alpha < 1$ if they both have points on a common edge $e \in E \in V(A)$, because E divides S_l and S_r. Vertex v is concave, which means that the angle between its corresponding edges $e_1, e_2 \in A$ is smaller than $180°$. e is the bisector of e_1 and e_2, which means that $\angle e e_1 = \angle e e_2 < 90°$. Therefore, points on E, will not

move along E to v. Instead they move to their closest point on e_1 or e_2, which is not v. Therefore, S_1 and S_2 will move in a different direction, whilst staying in their corresponding Voronoi cells, causing them to lose connectivity.

Points in both shapes move to their closest point in the other shape. This means that every point $p_a \in A$, only lies on the closest path from p_a to B, and every point $p_r \in B$, only lies on the shortest path from p_r to A. We know that the shortest paths from p_a to B and p_r to A do not have the same endpoints, because, the lemma states that no $p_a \in A$ has p_r as its closest point. Therefore, only one endpoint can be shared. Points are only located at the endpoints of the shortest paths when $\alpha = 0$ or $\alpha = 1$. We also know that two shortest paths between A and B that do not share both endpoints do not intersect at any other point than the endpoint. This is proven by de Kogel et al. in Lemma 3 of [15]. Therefore, p_r, can only overlap with a point from A when $\alpha = 1$ or $\alpha = 0$, depending on the direction of the morph. This means that no $p_r \in B$ overlaps with any point $p_a \in A$, for any $0 < \alpha < 1$. This completes the proof.

□

7.2 Reducing components

We propose a new morph, based on the Voronoi morph, that is able to reduce the number of components without the need to increase the surface area of intermediate shapes. We call the morph the Voronoi Component Reduction (VCR) morph. This morph identifies which slices converge to extra components in the Voronoi morph, and moves them along with neighbouring slices that are part of a bigger component if the Hausdorff distance allows it. Smaller components always move along with bigger components because visually, small components look as if they move away from the bigger component that represents the input shapes. A slice is considered a neighboring slice if it is directly connected to the slice in the input shape. Recall that a component is a disconnected part of the shape, and that components can be composed of multiple slices. In the Voronoi morph, all components have a *target*. This targets is the site of the Voronoi cell by which the slice is partitioned: a point, line segment or interior component of the other shape. The target of a slice can be set to that of a neighboring slice. In that case, the slice is *redirected*. We call redirected targets *secondary targets*. There are several types of valid secondary targets to which a slice can be redirected. A secondary target is valid when the Hausdorff distance between the slice and target is smaller than of the two input shape and it satisfies one of the following conditions:

- it is the target of a neighboring slice that is part of a bigger component;
- it is the target of a neighboring slice, part of another component, that has valid secondary target in a bigger component;
- it is the target of a neighboring slice in the same component that has a valid secondary target.

Whenever a slice has a valid secondary target, it will be redirected, making the secondary target the new target of the slice. If a slice has multiple valid secondary targets, it will be redirected to the secondary target, to which the directed Hausdorff distance is the closest. Assuming the targets final targets are given, we formally define our new morph U_α as:

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

where $v_c(a,t_s)$ the closest point on the final target t_s of the slice in which a is located. If a slice has no valid secondary targets, this is will be the closest on B. If it has secondary targets, the target is the closest secondary target based on the directed Hausdorff distance. Notice that if no

slice can be redirected, this results in the Voronoi morph. Secondary targets are only valid if their Hausdorff distance is smaller than that of the two input shapes. Therefore, we can prove U_α is a Hausdorff morph.

Theorem 4. *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 $d_H(A, U_\alpha) = \alpha$ and $d_H(B, U_\alpha) = 1 - \alpha$.*

Proof. Theorem 1 of [15] shows that $d_H(A, T_\alpha) = \alpha$ and $d_H(B, T_\alpha) = 1 - \alpha$. It proves this by showing that every point moves to their closest point in the other shape linearly. This results in all points whose distance to the other shape is smaller than $d_H(A, B)$, to not affect the $d_H(A, T_\alpha)$ and $d_H(B, T_\alpha)$, and point p whose distance to the other shape is exactly $d_H(A, B)$ to be $d_H(p, T_\alpha) = \alpha$ and $d_H(p, T_\alpha) = 1-\alpha$. Therefore, moving a point a on a straight line segment with an endpoint on a and an endpoint on the boundary of B, with a length smaller than $d_H(A, B)$, never affects $d_H(A, T_\alpha)$ and $d_H(B, T_\alpha)$. The same holds for a point b that is moved on a straight line segment with an endpoint on b and an endpoint on the boundary of A.

This morph only allows a slice S to be redirected to a target t, when $d_H(S, t) < d_H(A, B)$. Therefore, redirecting slices does not affect the Hausdorff distance of intermediate shapes. Thus, $d_H(A, T_\alpha) = d_H(A, U_\alpha)$ and $d_H(B, T_\alpha) = d_H(B, U_\alpha)$. Hence, the statement in the theorem follows. □

Note that despite the goal to reduce components, it is theoretically possible that this approach results in more components. In Figure 13 a potential component in the Voronoi morph, composed of three slices, of an intermediate shape is illustrated. If slice $S_{3\alpha}$ where to be redirected, while $S_{1\alpha}$ and $S_{2\alpha}$, can not be redirected due to the Hausdorff constraint, the intermediate shape possesses an extra component in the VCR morph.

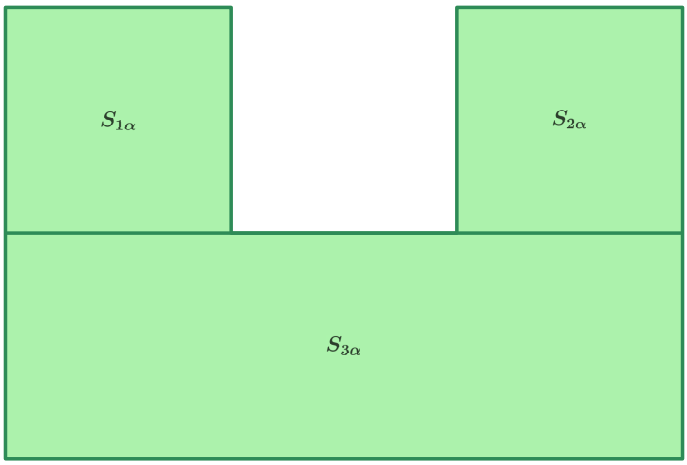

Figure 13: A potential component composed of three slices at a given $0 < \alpha < 1$.

7.3 Mixed variant

De Kogel et al. [15] show that the Voronoi morph introduces many slits into the boundary of intermediate shapes. VCR, has the same problem. They solve this problem by introducing a variant morph that *closes*, small gaps and holes, called the mixed morph. A mixed variant of the VCR morph can also be used to resolve these slits. The mixed VCR morph $U^m_{\alpha,\varphi}$, is defined as follows:

$$U^m_{\alpha,\varphi} := ((U_\alpha(A, B) \oplus D_\varphi) \ominus D_\varphi) \cap S_\alpha,$$

where $\ominus$ denotes the Minkowski difference, defined as $A \ominus B := (A^c \oplus B)^c$, where A^c is the complement of A. D_φ denotes a disc of radius φ. This means that $U^m_{\alpha,0} = U_\alpha$. To make sure it is a Hausdorff morph, the intersection with S_α is taken.

7.4 Algorithm

The algorithm to compute the VCR morph uses the algorithm to compute the Voronoi morph as the first step. To compute U_α we assume A and B are (sets of) polygons that may contain holes. The basic algorithm on these input sets works as follows:

1. Compute T_α at $\alpha = 1/2$ as described in De Kogel et al. [15]. This results in a set of slices, in which each slice has a target and is morphed halfway to their target in the other slice. A target can be a vertex, edge or interior component of the other input shape.

2. Determine which slices belong to which components in the halfway Voronoi morph for both input shapes separately. In our implementation, we dilate a slice by a tiny amount and check if it overlaps with other slices to determine if they belong to the same component.

3. For each slice we determine its neighbouring slices in the source shape. This is also implemented using a tiny dilation.

4. For each slice determine the shortest distance to a neighbouring component in the initial shape in terms of slices, using any graph searching algorithm such as breadth-first search. Two adjacent slices have a distance of 1. If two slices are both adjacent to slice s, but not adjacent to each other their distance is 2.

5. For both shapes A and B, sort all slices separately based on two criteria: (1) on the size of the component they belong to, from small to large, and (2) within each component, based on their slice distance to an adjacent component in the initial shape.

6. For each slice, in order of the previously sorted slices, determine which neighbours in the original shape are valid secondary targets. The criteria for a valid target are described in Section 7.2. If a slice has multiple valid targets, the closest valid secondary target in terms of the directed Hausdorff distance from the slice to the target, will be set as the new target. If no valid secondary targets exist, the old target remains. If a slice is redirected, the slice is marked as a follower of the slice that belongs to the new target, and every follower of a redirected slice has their target set to the new valid target if the Hausdorff distance allows for it. If the Hausdorff distance does not allow it, the initial primary target for that slice is set as the target.

7. Each slice in A is scaled and translated towards their target in B. If the target is an interior component of the other shape, the slices overlap and will be stationary throughout the morph. If the target is a vertex, the slice will be uniformly scaled towards that vertex by a fraction of α. If the target is an edge, the slice will scale perpendicular to the supporting line of that edge by a factor of α. Slices in B are handles the same, except that they are scaled by a factor of $1 - \alpha$.

8. Combine all slices of A and B together. In our implementation, this is implemented using a tiny dilation.

We sort components from small to large, in order to ensure that larger components can move along with a smaller component if that smaller component is redirected to a component that is

even larger. We also sort on the distance to neighbouring components in terms of slices to make sure slices with no direct neighbouring component in the initial shape can be redirected if a neighbouring slice that is closer to another component is redirected towards a new target.

$U^m_{\alpha,\varphi}$ can simply be computed by computing U_α and S_α, dilating and eroding U_α by a distance of φ, and intersecting the result with S_α.

7.5 Experiments

We compare the Voronoi, VCR, and both mixed variant morphs experimentally on the three data sets described in Section 4.2. For each experiment, the area, perimeter, total angular change, 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 variants is universally set to the function described in Equation (1), based on the test performed in Sections 6.1.1 and 6.2.1.

7.6 Results

Tables 5 and 6 show a summary of the surface area and perimeter measurements. Table 7 contains a summary of the number of components and holes for the animals data set; the other data sets are excluded because the inputs have different numbers of components. A full overview of the topological measurements for all experiments is shown in Tables 8 and 9 in Appendix B. Due to numerical precision issues, the morphs sometimes contain spurious holes (e.g. no intermediate shape in our experiment with the words try and it, should contain two holes in the Voronoi morph). Figure 14 displays that the VCR morph is a tiny bit worse in terms of area than the Voronoi morph. In the Voronoi morph, slices never overlap, because they all move to their points in the other shape. In the VCR morph, slices can overlap, because their directions can be changed to limit the number of components. This overlap seems to result in a lower surface area. The overlapping slices can also explain the fact surface area change over time is less smooth in the VCR morph. This is less noticeable in the countries data set. This could be a result of the countries data set having more test shapes, in which the average converges to a *smoother* curve. The mixed VCR morph, also has a lower surface area than the mixed morph, due to the same reason. Figure 14 shows that the mixed morph's surface area stays very close to one for all tested values of alpha on the countries data set.

In terms of perimeter, the VCR and mixed VCR morphs are both an improvement of their Voronoi and mixed morphs counterparts. The perimeter of the VCR morph is significantly lower than that of the Voronoi morph. The mixed VCR morph slows a slight improvement in terms of perimeter compared to the mixed morph. This effect can be explained by the reduction of components. If a slice disconnects from another slice, their previously shared boundary is added to the total perimeter.

The main purpose of our VCR morph is to reduce the number of extra components created in the Voronoi morph. It achieves this purpose well; the average number of components created on the animal data set in the VCR morph is more than three times less than that in the Voronoi morph. In Figure 17, we see that all extra components in the Voronoi shark-spider morph, are not present in the VCR morph. The tail of the shark stays intact for any intermediate shape in the VCR morph. The mixed and mixed VCR morphs further reduces the number components of their Voronoi and VCR counterparts. The mixed and mixed VCR morphs do however create some undesirable outcomes in the countries data set. Components that are disconnected in one of the input shapes can be merged as a result of the closing operator. This is illustrated in Figure 19: the islands on the bottom left of Spain partially merged in both mixed variants.

	Voronoi		Mixed		VCR		Mixed VCR	
Category	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Animals	0.976	0.022	0.999	0.017	0.964	0.076	0.983	0.076
Countries	0.968	0.036	1.008	0.047	0.963	0.061	0.994	0.062
Text	1.027	0.031	1.035	0.030	1.021	0.041	1.027	0.038

Table 5: The distributions of normalized surface areas for each morphing method over all experiments for all nine tested values of α, separated by experiment category.

	Voronoi		Mixed		VCR		Mixed VCR	
Category	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Animals	1.531	0.322	1.124	0.123	1.408	0.278	1.058	0.117
Countries	1.528	0.415	1.060	0.149	1.397	0.302	1.019	0.118
Text	1.236	0.286	1.130	0.177	1.161	0.191	1.082	0.116

Table 6: The distributions of normalized perimeters for each morphing method over all experiments for all nine tested values of α, separated by experiment category.

	Voronoi		Mixed		VCR		Mixed VCR	
Category	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Components	11.262	4.686	4.210	2.480	3.710	2.263	3.016	2.024
Holes	0.282	0.589	0.337	0.663	1.690	1.787	0.508	0.770

Table 7: The distributions of the number of components and holes for each morphing method for all tested values of α except 0 and 1. Only the animal data set is included, as these shapes only have one component and no holes.

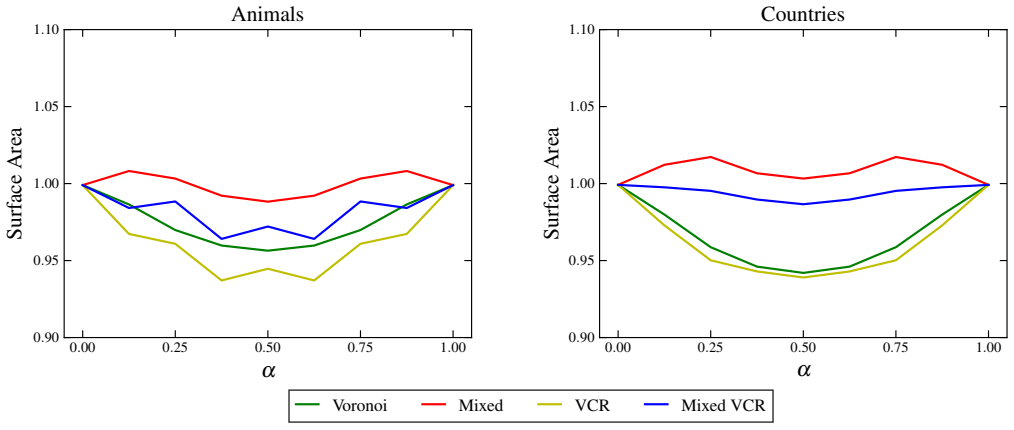

Figure 14: The normalized average surface area over all experiments as a function of α, for both the animals and countries data sets.

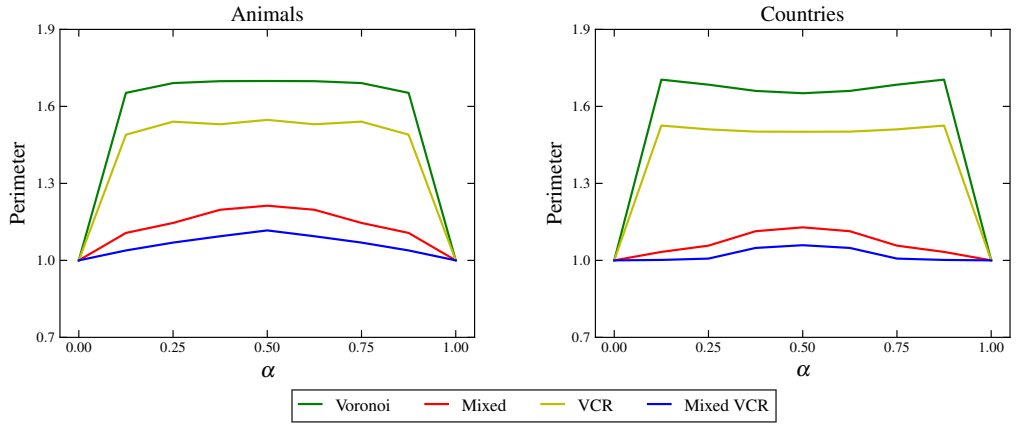

Figure 15: The normalized average perimeter over all experiments as a function of α, for both the animals and countries data sets.

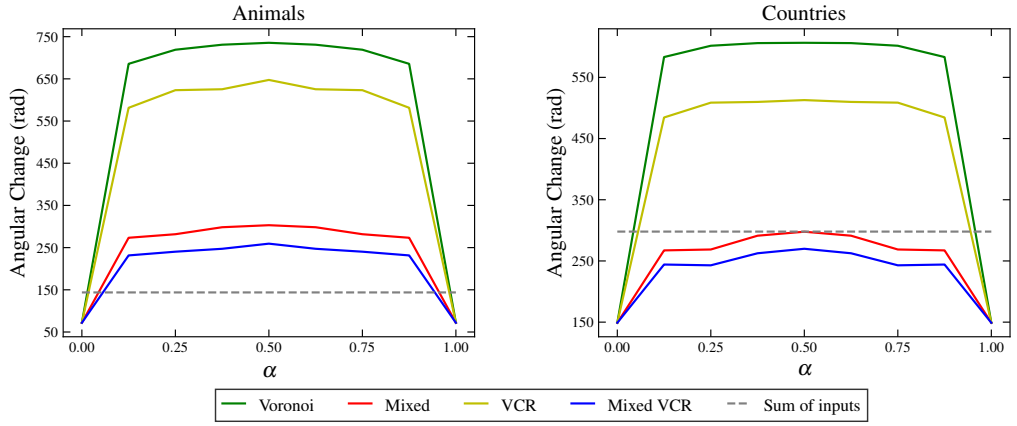

Figure 16: The average total angular change over all experiments as a function of α, for both the animals and countries data sets.

The total angular change on our VCR and mixed VCR morphs are lower than their respective Voronoi and mixed morphs. Each component adds at least 2π to the total angular change. Therefore, reducing the number of components leads to a lower total angular change. The shape itself does not seem to differ too much between our VCR and mixed VCR morphs, and their respective Voronoi and mixed morphs. The difference in the component count is larger between the Voronoi and VCR morphs, than that of the mixed variants. Therefore, the difference in total angular change between the Voronoi and VCR morphs is larger than that of the mixed and mixed VCR morphs.

The change in topology is however not improved in terms of the number of holes. Our VCR morph's intermediate shapes contain more holes than intermediate shapes in the Voronoi morph. Figure 20 displays how slits in the Voronoi morph, can result in holes in our VCR morph, if a component is redirected in such a way that it overlaps with the opening of the slit. Our Mixed VRC morph manages to fill some of these holes, resulting in a significant lower hole count on the animals data set. It does however in some cases also create holes that were not present in the VCR morph. Figure 17 shows the mixed and Mixed VCR morphs, containing holes that are not present in the Voronoi and VCR morphs.

Visually (Figures 17-21, party in Appendix A), our VCR morph seems to better conserve

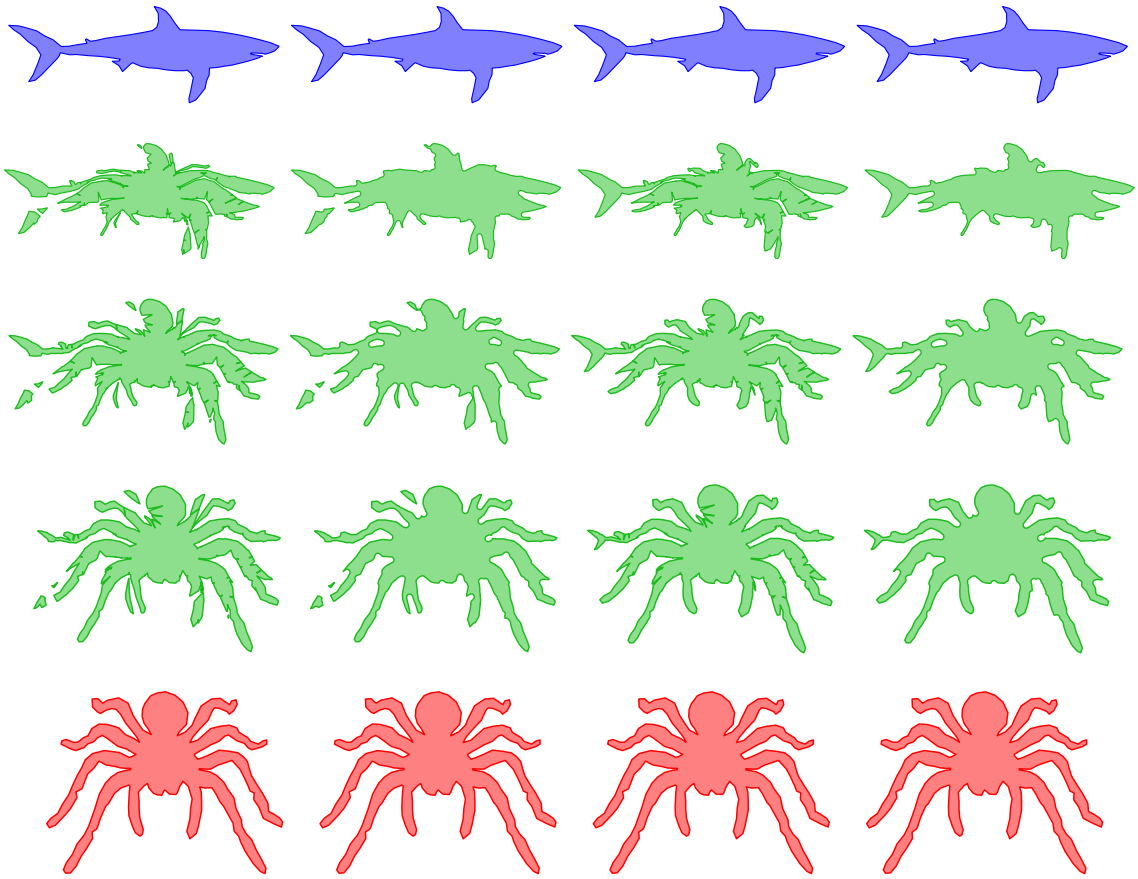

Figure 17: Intermediate shapes for $\alpha \in \{0, 1/4, 1/2, 3/4, 1\}$ when morphing between outlines of a shark and spider. The columns show the Voronoi morph, mixed morph, VCR morph and mixed VCR morph from left to right.

subtle details from the input shapes. The shark's tail is kept intact in all intermediate shapes when morphing between the shark and spider. In the Voronoi morph, half of the tail splits away from the main component, resulting in two extra components that do not reps resent any details in the input shapes. Even in cases where extra components are still created, lowering the number of components increases the visual quality of the morph. The hind leg of the ostrich, is more recognisable in the VCR morph than in the Voronoi morph, where it is split into more components (Figure 18 in Appendix A).

Morphing between words is a challenge for the Voronoi morph. Our VCR morph better conserves some of the structure of certain letters. Figure 21 illustrates that the tittle of the letter i does not split into multiple components, and is therefore more recognizable. The letter y also distorts less in the VCR morph. Regardless of these improvements, at $\alpha = 1/2$ it is still hard to recognize the input words due to the fact that words are comprised of individual letters that are all being merged together, as distinct letters are not taken into account.

8 Conclusion

We delivered three major contributions: the total angular change metric, experimental tests on the mixed morph with different parameter settings, and the Voronoi component reduction (VCR) Hausdorff morph and its mixed variant.

The total angular change over time is a quantitative metric that is experimentally demonstrated to effectively distinguish different types of morphs in a predictable manner. By iteratively removing details using a larger dilation parameter on the mixed morph, the angular change accurately reflects the loss of concave detail. This metric can effectively explain visual results. It does however not necessarily indicate what happens in a morph without looking at the morph itself, because a higher total angular change can be explained by two shape features: concave details and extra components. Both features, usually result in a worse visual morph, therefore it can still be used to accurately determine which of two outputs is visually better in most cases in the tested data sets.

To improve the quality of the existing mixed morph, we experimentally tested the morph with different parameter settings. We showed that changing the dilation parameter has a big impact on the quality of the morph's output. Usually, when the dilation parameter is changed to improve one recorded feature such as the number of components, this results in other features, such as the surface area, to suffer. We found changing the radius of the disc for the closing operator over time is the best way to balance the quality of all features. When α is close to 0 or 1, the disc's radius can be small, to not affect the perimeter too much, while still keeping the number of components low by gradually increasing the radius of the closing disc.

Even with the right balance, the mixed morph still suffers from the extra components created in the Voronoi morph. It might merge some components using the closing operator, but that also ussually results in details that are not present in the input shapes. Our VCR morph solves this issue and still satisfies the Hausdorff property. We have shown experimentally that the number of components in the VCR morph is significantly lower than that of the Voronoi and mixed morphs. We also demonstrated that the VCR morph keeps more visual details of the input shapes in tact during the morphs. This results in intermediate shapes being more recognizable in the VCR morph than those formed by the Voronoi and mixed morphs. To solve the issue the Voronoi and VCR morphs have regarding their perimeter, caused by slits, we also introduced a mixed variant of the VCR morph and experimentally demonstrate that this mixed variant is able to effectively reduce the perimeter caused by slits. This leaves a morph that better preserves details of both input shapes without adding too much extra details.

It must be noted that the VCR morph loses some of the conceptual simplicity of the Voronoi and mixed morphs. Selecting correct targets that reduce the number of components significantly increases the complexity of the algorithm. It does however leave some elegant simplicity from the way slices only move along with direct neighbours in their source shape, resulting in the morph working on inputs that contain multiple components without the need to adjust the algorithm or implementation.

9 Future work

There are several ways the Voronoi and VCR morphs might be further improved. Some open questions that could lead to further research will be briefly discussed here.

The VCR morph might be improved when the criteria on choosing the best valid secondary target change, because different targets cause the slice to deform in different ways. Currently, if multiple valid secondary targets exist, the closest in terms of the directed Hausdorff distance is

selected. The best secondary target could be selected in several other manners. The first could be to move along with the neighbouring slice that has the largest shared boundary. Another approach could be to use a different distance measurement. The directed Hausdorff distance measures the distance from the furthest point in the slice to the target. We might be able to more accurately predict the amount of deformation if we take the distances of all vertices to the target into account.

In this thesis, we briefly tested the capabilities that changing the radius for the closing operator has. It showed promising, but limited results. Only one variation was tested. It is therefore likely that further experimental research could lead to improvement on the mixed variants. Different functions can be tested and versions that adaptively change to balance all, or some, recorded features might prove to significantly improve the quality of these morphs for all values of α.

Our implementation normalizes the position of input shapes to have a common centroid. Different ways of positioning the input shapes might lead to significant improvements on the Voronoi, VCR, and mixed variant morphs. The shapes' position can for example be normalized by minimizing the Hausdorff distance, or by maximizing the overlap between the input shapes.

The Voronoi, VCR and mixed variant morphs are created to morph between 2D shapes. Conceptually they are quite simple to adjust to work on 3D shapes. Most metrics that determine the quality of the morph can also easily be used when morphing between 3D shapes. The surface area can be substituted for the volume, and the perimeter can be substituted for the surface area. The total angular change can not easily be substituted for a 3D version, since 3D shapes do not allow for a *walk* around a given boundary loop. Other metrics such as compactness or eccentricity might turn out to be accurate indicators of a morph's quality that work on both 2D and 3D shapes.

References

[1] Alexandra Albu, Trevor Beugeling, and Denis Laurendeau. “A Morphology-Based Approach for Interslice Interpolation of Anatomical Slices From Volumetric Images”. In: *IEEE transactions on bio-medical engineering* 55 (Aug. 2008), pp. 2022–38.

[2] Marc Alexa, Daniel Cohen-Or, and David Levin. “As-Rigid-as-Possible Shape Interpolation”. In: *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’00. USA: ACM Press/Addison-Wesley Publishing Co., 2000, pp. 157–164. ISBN: 1581132085.

[3] Boris Aronov, Raimund Seidel, and Diane Souvaine. “On compatible triangulations of simple polygons”. In: *Computational Geometry* 3.1 (1993), pp. 27–35.

[4] Gill Barequet, Michael T Goodrich, Aya Levi-Steiner, and Dvir Steiner. “Contour interpolation by straight skeletons”. In: *Graphical Models* 66.4 (2004), pp. 245–260.

[5] Gill Barequet and Micha Sharir. “Piecewise-Linear Interpolation between Polygonal Slices”. In: *Computer Vision and Image Understanding* 63.2 (1996), pp. 251–272.

[6] Gill Barequet and Amir Vaxman. “Nonlinear Interpolation between Slices.” In: *International Journal of Shape Modeling* 14 (Jan. 2008), pp. 39–60.

[7] Gill Barequet and Amir Vaxman. “Reconstruction of Multi-Label Domains from Partial Planar Cross-Sections”. In: *Comput. Graph. Forum* 28 (July 2009), pp. 1327–1337.

[8] Quirijn W. Bouts, Irina Kostitsyna, Marc J. van Kreveld, Wouter Meulemans, Willem Sonke, and Kevin Verbeek. “Mapping polygons to the grid with small Hausdorff and Fréchet distance”. In: *Proceedings of the 24th Annual European Symposium on Algorithms* (2016).

[9] K.A. Buchin, E.W. Chambers, T.A.E. Ophelders, and B. Speckmann. “Fréchet isotopies to monotone curves”. English. In: 33rd European Workshop on Computational Geometry (EuroCG 2017), EuroCG 2017 ; Conference date: 05-04-2017 Through 07-04-2017. Apr. 2017, pp. 41–44.

[10] Kevin Buchin, Maike Buchin, Wouter Meulemans, and Bettina Speckmann. “Locally Correct Frechet Matchings”. In: *CoRR* abs/1206.6257 (2012). arXiv: 1206.6257.

[11] T. Chan and Wei Zhu. “Level set based shape prior segmentation”. In: *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*. Vol. 2. 2005, 1164–1170 vol. 2.

[12] Daniel Cohen-Or, Amira Solomovic, and David Levin. “Three-Dimensional Distance Field Metamorphosis”. In: 17.2 (1998).

[13] Craig Gotsman and Vitaly Surazhsky. “Guaranteed intersection-free polygon morphing”. In: *Computers & Graphics* 25.1 (2001). Shape Blending, pp. 67–75.

[14] Adobe Systems Incorporated. “Designing Multiple Master Typefaces”. In: (July 2015).

[15] Lex de Kogel, Marc van Kreveld, and Jordi L. Vermeulen. “Abstract Morphing Using the Hausdorff Distance and Voronoi Diagrams”. In: *30th Annual European Symposium on Algorithms (ESA 2022)*. Vol. 244. Leibniz International Proceedings in Informatics (LIPIcs). 2022, 74:1–74:16. ISBN: 978-3-95977-247-1.

[16] Marc van Kreveld, Tillmann Miltzow, Tim Ophelders, Willem Sonke, and Jordi L. Vermeulen. “Between shapes, using the Hausdorff distance”. In: *Computational Geometry* 100 (2022), p. 101817.

[17] Anjum I. Rohra and Ramesh K. Kulkarni. “Survey on Recent trends in Image Morphing Techniques”. In: *2019 International Conference on Smart Systems and Inventive Technology (ICSSIT)*. 2019, pp. 19–23.

[18] Justin Solomon, Fernando Goes, Gabriel Peyré, Marco Cuturi, Adrian Butscher, Andy Nguyen, Tao Du, and Leonidas Guibas. “Convolutional Wasserstein Distances: Efficient Optimal Transportation on Geometric Domains”. In: *ACM Transactions on Graphics* 34 (July 2015).

A More example morphs

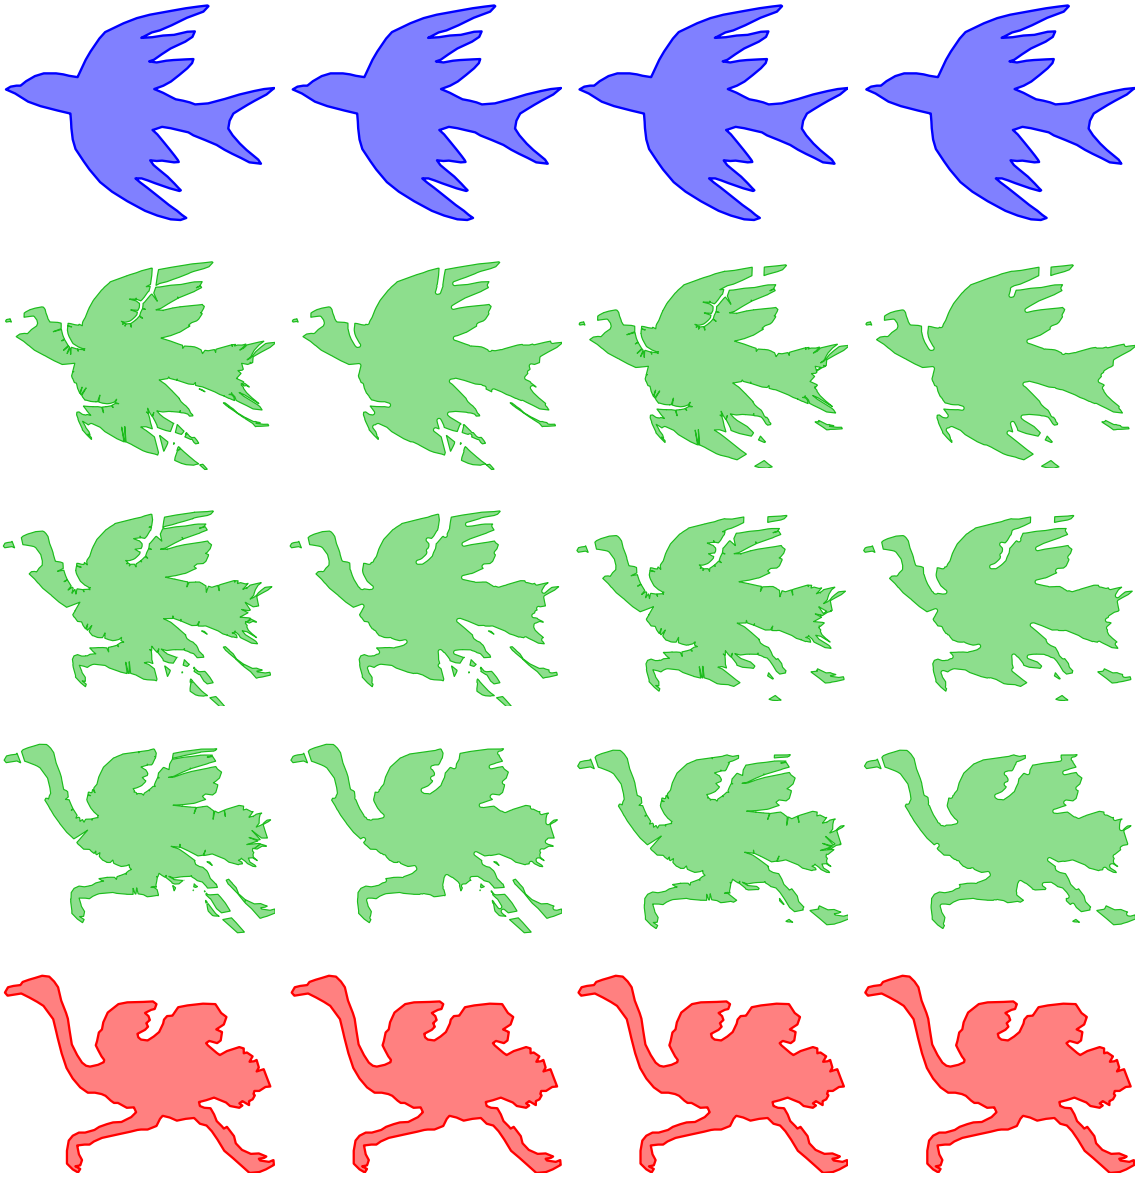

Figure 18: Intermediate shapes for $\alpha \in \{0, 1/4, 1/2, 3/4, 1\}$ when morphing between outlines of a bird and ostrich. The columns show the Voronoi morph, mixed morph, VCR morph and mixed VCR morph from left to right.

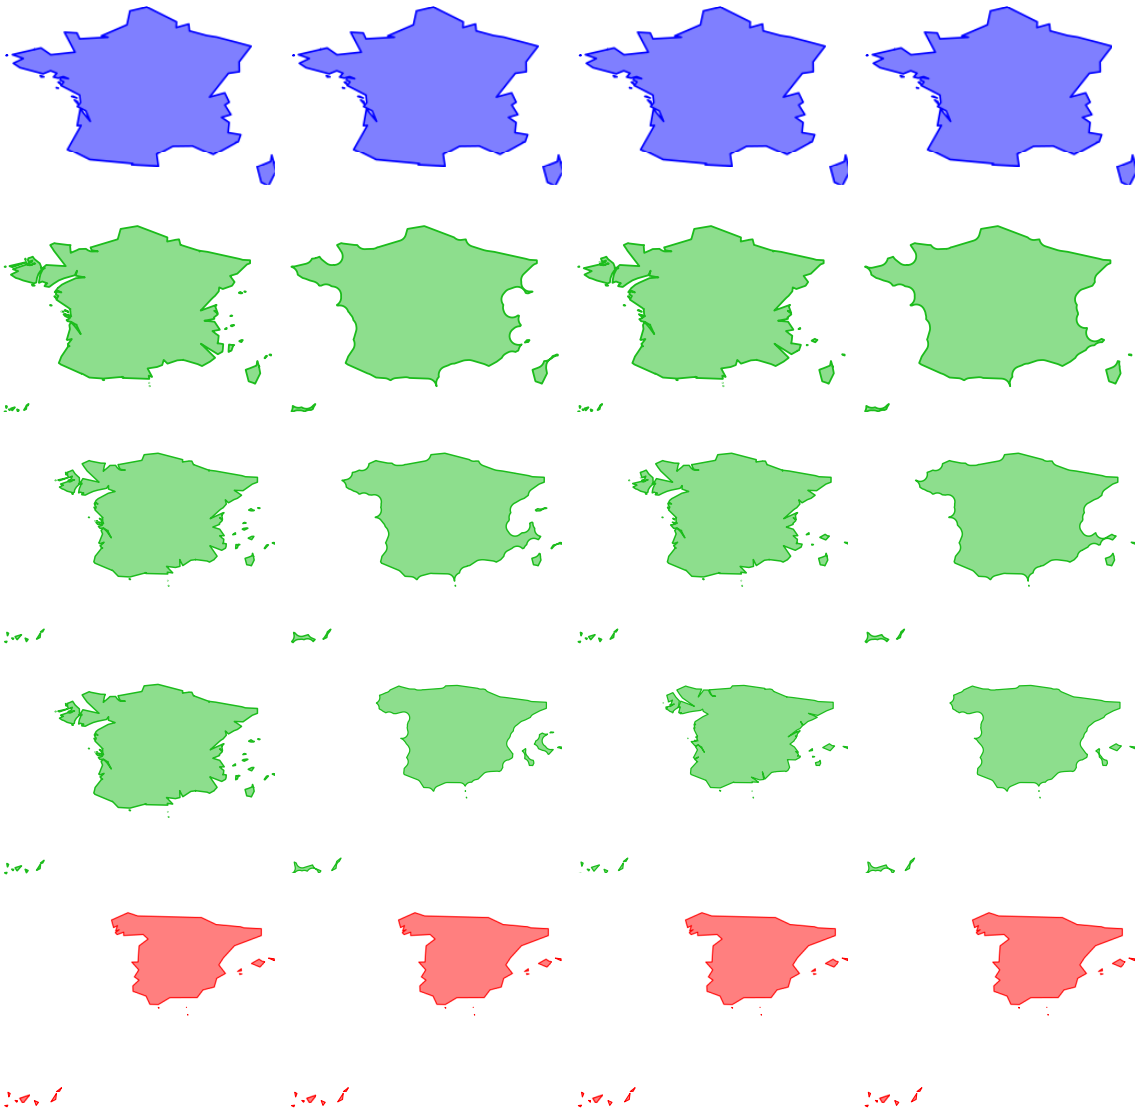

Figure 19: Intermediate shapes for $\alpha \in \{0, 1/4, 1/2, 3/4, 1\}$ when morphing between outlines of France and Spain. The columns show the Voronoi morph, mixed morph, VCR morph and mixed VCR morph from left to right.

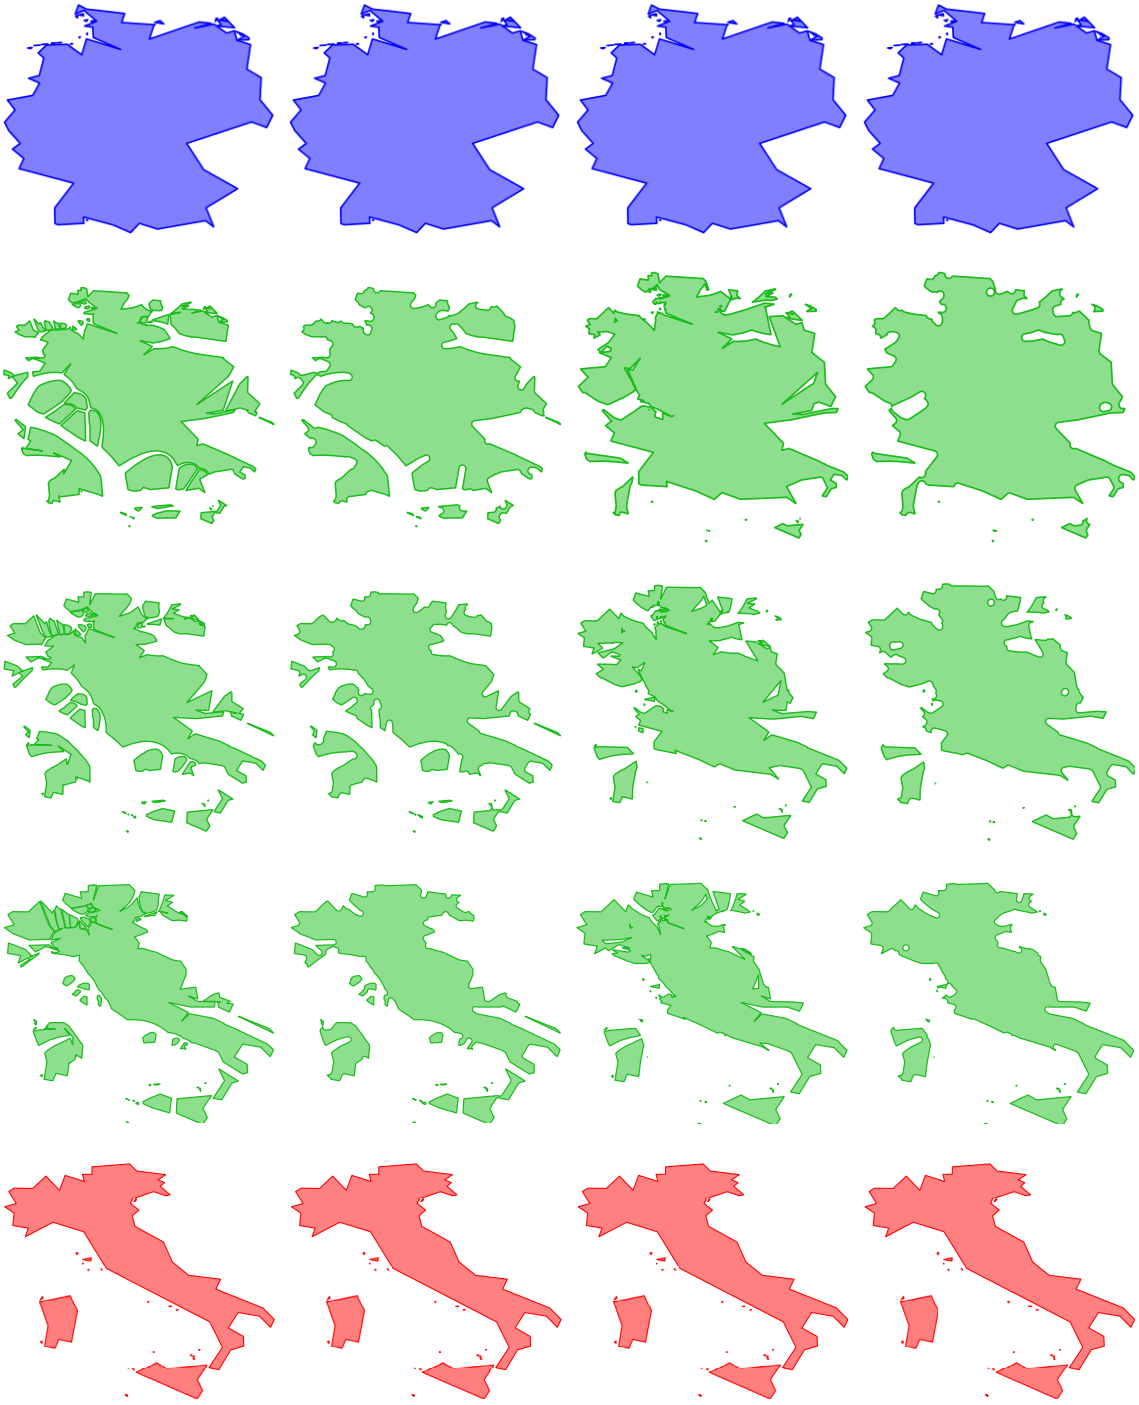

Figure 20: Intermediate shapes for $\alpha \in \{0, 1/4, 1/2, 3/4, 1\}$ when morphing between outlines of Germany and Italy. The columns show the Voronoi morph, mixed morph, VCR morph and mixed VCR morph from left to right.

Figure 21: Intermediate shapes for $\alpha \in \{0, 1/4, 1/2, 3/4, 1\}$ when morphing between outlines of of the words try and it. The columns show the Voronoi morph, mixed morph, VCR morph and mixed VCR morph from left to right.

B Topology tables

Table 8: The minimum and maximum number of components and the maximum number of holes for each experiment, separated by the Voronoi and VCR morph types.

	Voronoi			VCR		
Experiment	min	max	holes	min	max	holes
bird → butterfly	1	8	0	1	4	1
bird → cat	1	10	0	1	1	6
bird → dog	1	10	1	1	4	2
bird → horse	1	17	1	1	3	5
bird → ostrich	1	16	0	1	6	2
bird → shark	1	11	0	1	2	2
bird → spider	1	15	2	1	2	3
bird → turtle	1	14	0	1	5	0
butterfly → cat	1	3	0	1	1	0
butterfly → dog	1	5	1	1	2	2
butterfly → horse	1	21	1	1	8	3
butterfly → ostrich	1	6	0	1	4	6
butterfly → shark	1	3	0	1	2	0
butterfly → spider	1	17	1	1	5	4
butterfly → turtle	1	9	1	1	2	1
cat → dog	1	9	0	1	5	3
cat → horse	1	9	0	1	3	2
cat → ostrich	1	6	0	1	2	0
cat → shark	1	6	1	1	2	1
cat → spider	1	12	2	1	4	4
cat → turtle	1	10	1	1	5	1
dog → horse	1	13	2	1	7	3
dog → ostrich	1	12	2	1	4	3
dog → shark	1	12	0	1	1	0
dog → spider	1	19	1	1	7	8
dog → turtle	1	12	0	1	4	1
horse → ostrich	1	21	0	1	9	3
horse → shark	1	13	1	1	3	2
horse → spider	1	17	3	1	6	9
horse → turtle	1	17	1	1	7	4
ostrich → shark	1	11	1	1	4	6
ostrich → spider	1	22	3	1	15	4
ostrich → turtle	1	11	2	1	2	0
shark → spider	1	13	0	1	1	1
shark → turtle	1	8	0	1	2	0
spider → turtle	1	15	1	1	4	4
austria → belgium	1	1	0	1	1	0
austria → croatia	1	20	2	1	19	5
austria → czechia	1	2	0	1	1	0

Table 8: (continued from last page)

	Voronoi			VCR		
Experiment	min	max	holes	min	max	holes
austria → france	1	14	1	1	10	5
austria → germany	1	20	1	1	20	1
austria → greece	1	71	1	1	68	9
austria → ireland	1	9	0	1	5	0
austria → italy	1	29	0	1	25	4
austria → netherlands	1	11	4	1	9	4
austria → poland	1	2	0	1	2	0
austria → spain	1	20	0	1	15	1
austria → sweden	1	21	7	1	19	8
belgium → croatia	1	22	0	1	19	1
belgium → czechia	1	1	0	1	1	0
belgium → france	1	11	2	1	10	10
belgium → germany	1	21	2	1	20	3
belgium → greece	1	73	1	1	76	7
belgium → ireland	1	6	0	1	6	0
belgium → italy	1	24	0	1	24	5
belgium → netherlands	1	11	3	1	9	10
belgium → poland	1	1	0	1	1	0
belgium → spain	1	17	0	1	16	0
belgium → sweden	1	22	9	1	19	11
croatia → czechia	1	23	1	1	19	3
croatia → france	10	43	1	10	26	8
croatia → germany	19	52	2	19	35	6
croatia → greece	19	111	0	19	79	6
croatia → ireland	4	33	0	4	21	5
croatia → italy	19	59	0	19	40	6
croatia → netherlands	9	31	3	9	30	7
croatia → poland	1	29	1	1	19	3
croatia → spain	15	49	0	15	32	0
croatia → sweden	19	53	7	19	30	7
czechia → france	1	11	1	1	10	8
czechia → germany	1	20	2	1	20	2
czechia → greece	1	76	1	1	69	12
czechia → ireland	1	6	0	1	4	4
czechia → italy	1	24	0	1	23	4
czechia → netherlands	1	11	2	1	9	7
czechia → poland	1	2	0	1	1	0
czechia → spain	1	16	0	1	16	1
czechia → sweden	1	23	9	1	19	7
france → germany	10	32	1	10	21	7
france → greece	10	90	2	10	73	9
france → ireland	4	16	1	4	11	4
france → italy	10	43	1	10	32	7
france → netherlands	9	20	4	9	14	9

Table 8: (continued from last page)

	Voronoi			VCR		
Experiment	min	max	holes	min	max	holes
france → poland	1	10	1	1	10	9
france → spain	10	32	2	10	24	4
france → sweden	10	32	6	10	22	14
germany → greece	20	88	8	20	68	20
germany → ireland	4	24	2	4	20	5
germany → italy	20	51	2	20	27	13
germany → netherlands	9	35	4	9	25	8
germany → poland	1	20	1	1	21	5
germany → spain	15	36	2	15	30	6
germany → sweden	19	41	8	19	24	15
greece → ireland	4	76	1	4	72	10
greece → italy	22	112	1	22	84	11
greece → netherlands	9	81	2	9	83	21
greece → poland	1	76	2	1	71	16
greece → spain	15	101	2	15	79	4
greece → sweden	19	112	7	19	73	12
ireland → italy	4	34	0	4	25	4
ireland → netherlands	4	17	2	4	12	6
ireland → poland	1	7	1	1	5	2
ireland → spain	4	24	0	4	18	2
ireland → sweden	4	25	6	4	19	13
italy → netherlands	9	36	2	9	29	6
italy → poland	1	25	0	1	23	3
italy → spain	15	47	0	15	42	4
italy → sweden	19	53	10	19	36	11
netherlands → poland	1	11	3	1	9	4
netherlands → spain	9	24	3	9	21	4
netherlands → sweden	9	32	7	9	23	14
poland → spain	1	16	0	1	15	0
poland → sweden	1	21	10	1	19	15
spain → sweden	15	39	8	15	28	16
wish → luck	4	22	2	4	10	8
kick → stuff	5	20	1	5	11	5
try → it	3	13	2	3	7	4
f serif → f sans	1	1	0	1	1	0
i serif → i sans	2	2	0	2	2	0
u serif → u sans	1	1	0	1	1	0

Table 9: The minimum and maximum number of components and the maximum number of holes for each experiment, separated by the mixed and mixed VCR morph types.

	Mixed			Mixed VCR		
Experiment	min	max	holes	min	max	holes
bird → butterfly	1	4	2	1	4	2
bird → cat	1	4	0	1	1	1
bird → dog	1	5	1	1	3	1
bird → horse	1	5	1	1	3	3
bird → ostrich	1	10	1	1	6	0
bird → shark	1	5	0	1	2	0
bird → spider	1	6	2	1	2	1
bird → turtle	1	6	0	1	5	1
butterfly → cat	1	3	1	1	1	1
butterfly → dog	1	3	1	1	2	1
butterfly → horse	1	5	1	1	7	1
butterfly → ostrich	1	3	2	1	3	2
butterfly → shark	1	3	1	1	2	1
butterfly → spider	1	5	2	1	4	2
butterfly → turtle	1	3	1	1	2	2
cat → dog	1	1	1	1	3	2
cat → horse	1	5	1	1	3	1
cat → ostrich	1	3	0	1	2	0
cat → shark	1	4	0	1	2	0
cat → spider	1	7	2	1	4	2
cat → turtle	1	5	1	1	5	0
dog → horse	1	4	1	1	6	2
dog → ostrich	1	5	1	1	3	1
dog → shark	1	5	1	1	1	0
dog → spider	1	12	2	1	6	3
dog → turtle	1	4	0	1	3	1
horse → ostrich	1	12	1	1	8	3
horse → shark	1	6	1	1	2	1
horse → spider	1	7	3	1	5	4
horse → turtle	1	8	0	1	7	1
ostrich → shark	1	5	2	1	4	1
ostrich → spider	1	15	1	1	15	2
ostrich → turtle	1	3	0	1	1	0
shark → spider	1	4	3	1	1	2
shark → turtle	1	2	1	1	2	0
spider → turtle	1	7	5	1	4	4
austria → belgium	1	1	0	1	1	0
austria → croatia	1	19	4	1	19	4
austria → czechia	1	1	0	1	1	0
austria → france	1	10	0	1	10	0
austria → germany	1	20	0	1	20	0
austria → greece	1	68	4	1	68	3

Table 9: (continued from last page)

	Mixed			Mixed VCR		
Experiment	min	max	holes	min	max	holes
austria → ireland	1	4	0	1	4	1
austria → italy	1	22	1	1	22	1
austria → netherlands	1	9	3	1	9	3
austria → poland	1	2	0	1	2	0
austria → spain	1	15	1	1	15	0
austria → sweden	1	19	0	1	19	0
belgium → croatia	1	19	4	1	19	3
belgium → czechia	1	1	0	1	1	0
belgium → france	1	10	1	1	10	1
belgium → germany	1	20	2	1	20	2
belgium → greece	1	68	2	1	68	4
belgium → ireland	1	4	1	1	4	1
belgium → italy	1	22	1	1	22	1
belgium → netherlands	1	9	2	1	9	2
belgium → poland	1	1	1	1	1	1
belgium → spain	1	15	0	1	15	0
belgium → sweden	1	19	1	1	19	1
croatia → czechia	1	19	2	1	19	2
croatia → france	7	19	2	5	19	4
croatia → germany	4	20	2	2	20	3
croatia → greece	19	68	5	19	68	4
croatia → ireland	2	19	4	2	19	3
croatia → italy	13	22	1	19	23	2
croatia → netherlands	6	19	4	4	19	4
croatia → poland	1	19	2	1	19	3
croatia → spain	6	19	0	3	19	1
croatia → sweden	10	19	3	6	19	3
czechia → france	1	10	0	1	10	0
czechia → germany	1	20	2	1	20	2
czechia → greece	1	68	3	1	68	4
czechia → ireland	1	4	2	1	4	3
czechia → italy	1	22	0	1	22	1
czechia → netherlands	1	9	2	1	9	3
czechia → poland	1	1	0	1	1	0
czechia → spain	1	15	0	1	15	0
czechia → sweden	1	19	1	1	19	2
france → germany	3	20	2	3	20	1
france → greece	10	68	3	9	68	3
france → ireland	3	10	1	4	10	2
france → italy	10	22	1	10	22	2
france → netherlands	5	10	2	4	10	3
france → poland	1	10	2	1	10	2
france → spain	4	15	0	4	15	0
france → sweden	7	19	1	5	19	1

Table 9: (continued from last page)

	Mixed			Mixed VCR		
Experiment	min	max	holes	min	max	holes
germany → greece	6	68	1	7	68	1
germany → ireland	3	20	2	1	20	3
germany → italy	9	22	3	9	22	3
germany → netherlands	6	20	3	5	20	4
germany → poland	1	20	3	1	20	4
germany → spain	2	20	0	2	20	0
germany → sweden	9	20	4	5	20	4
greece → ireland	3	68	4	4	68	4
greece → italy	22	68	2	22	68	7
greece → netherlands	6	68	3	9	68	4
greece → poland	1	68	3	1	68	3
greece → spain	11	68	2	11	68	1
greece → sweden	9	68	2	8	68	2
ireland → italy	4	22	1	4	22	1
ireland → netherlands	3	9	3	2	9	5
ireland → poland	1	4	2	1	4	1
ireland → spain	2	15	0	3	15	0
ireland → sweden	3	19	2	2	19	2
italy → netherlands	9	22	2	8	22	3
italy → poland	1	22	1	1	22	1
italy → spain	5	22	0	9	22	0
italy → sweden	10	22	1	5	22	2
netherlands → poland	1	9	2	1	9	2
netherlands → spain	3	15	1	3	15	1
netherlands → sweden	5	19	2	5	19	3
poland → spain	1	15	0	1	15	0
poland → sweden	1	19	1	1	19	1
spain → sweden	4	19	1	4	19	0
wish → luck	4	17	0	4	9	1
kick → stuff	5	17	0	5	10	1
try → it	3	9	1	3	6	1
f serif → f sans	1	1	0	1	1	0
i serif → i sans	2	2	0	2	2	0
u serif → u sans	1	1	0	1	1	0