

©2016

Nathan Destler

ALL RIGHTS RESERVED

SENSITIVITY TO SHAPE DIFFERENCES ALONG MORPH-SPACES

by

NATHAN ROBERT JOHNSON DESTLER

A thesis submitted to the

Graduate School-New Brunswick

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Master of Science

Graduate Program in Psychology

Written under the direction of

Jacob Feldman

And approved by

---

---

---

New Brunswick, New Jersey

OCTOBER, 2016

## ABSTRACT OF THE THESIS

### Sensitivity to Shape Differences Along Morph-Spaces

By NATHAN ROBERT JOHNSON DESTLER

Thesis Director:

Jacob Feldman

We investigated the dimensions defining mental shape space, by measuring shape discrimination thresholds along "morph-spaces" defined by pairs of shapes. Given any two shapes, one can construct a morph-space by taking weighted averages of their boundary vertices (after normalization), creating a continuum of shapes ranging from the first shape to the second. Previous studies of morphs between highly familiar shape categories (e.g. truck and turkey) have shown elevated discrimination at the category boundary, reflecting a kind of "categorical perception" in shape space. However, these findings were restricted to known object shapes. Here, we use this technique to explore implicit categorical boundaries in spaces of unfamiliar shapes, where categories are defined not by familiar named types, but by the underlying "generative" structure of mental shape space. We further explore how probabilistic skeletal models of shape may explain discrimination and categorization of these unfamiliar shapes. In this study, subjects were shown two shapes at nearby points along a morph-space, and asked to judge whether they were the same or different, with an adaptive procedure used to

estimate discrimination thresholds at each point along the morph-space. We targeted several potentially important categorical distinctions, such one- vs. two-part shapes, two- vs. three-part shapes, changes in symmetry structure, and other "qualitative" distinctions. The results show strong consistency between subjects. Sensitivity ( $1/\text{difference threshold}$ ) is predicted by using a Bayesian probabilistic skeletal model to compute the probability of the standard shape being generated by the comparison shape's generative model, and vice versa. The results show that discrimination thresholds are not uniform over shape spaces. Instead, the results are consistent with the model, suggesting that a probabilistic generative framework drives shape discrimination.

## TABLE OF CONTENTS

Abstract.....	ii
List of Tables .....	v
List of Illustrations.....	vi
Introduction .....	1
Design of the Experiment .....	8
Participants .....	11
Experiments .....	12
Results.....	14
Discussion .....	17
Conclusion.....	27
Works Cited.....	1

## LIST OF TABLES

Table 1. Listing of Morph-Spaces: Example shapes and corresponding labels .....	13
Table 2. Experimental Conditions by Morph-Space and Normalization Method.....	14
Table 3. AIC Results by Shape and Normalization Condition .....	25

## LIST OF ILLUSTRATIONS

Figure 1. Shapes extracted from morph-spaces .....	4
Figure 2. Sensitivity curves for each one-part/two-part condition .....	16
Figure 3. Sensitivity curves for the first six two-part/three-part conditions .....	18
Figure 4. Sensitivity curves for the last two-part/three-part condition .....	19
Figure 5. Sensitivity curves for each ellipse-based condition .....	20
Figure 6. Sample shape skeletons.....	21

## INTRODUCTION

There are many ways to parameterize shapes. Some parameterizations are suggested by the computer science literature, such as computing and comparing image moments or applying deep-learning algorithms (Chellappa, 2016). However, while these techniques can be quite useful, they do not generally provide a good model for how human beings parameterize shapes. Humans have long been believed to sort shapes according to a categorical structure, with each category representing some semantically meaningful set of parameters (Rosch, 1973, Shepard, Hovland, & Jenkins, 1961). At the highest level, a chair might be said to have the parameters “can be sat in” and “contains two connected surfaces roughly perpendicular to one another.” At a much lower level, a circle might have the parameters “round” and “has an eccentricity of 0.” A shape’s category, then, is defined according to its parameter values, whether they be highly abstract (“can be sat in”) or mathematically concrete (“has an eccentricity of 0”).

Each parameter can be represented as a dimension in a high-dimensional category space (Ashby & Maddox, 1990). Scores on a given parameter can be assigned according to mathematical principles, or according to how strongly a given shape matches a certain semantic property. For example, in the “can be sat in” dimension, a chairs will tend to have high scores, while coat hangers will tend to have low scores, and rocks may be anywhere in the range, depending on the rock. In the “round” dimension, however, chairs may have both high and low scores, because a chair is not strongly defined by its roundness. We can see from these examples that the category “chair” consists of shapes with a narrow range of values for some parameters, and a broader



range of values for other parameters. Each given chair is a single point in the high-dimensional category space, and the chair category can be considered to occupy a volume in that space that contains all (or most) such points. We can define categories, then, as functions over the volumes they occupy in the category space (Ashby & Alfonso-Reese, 1995).


However, none of that really answers the question of how humans parameterize shapes. What defines the category space? How is it structured? How can we determine the shape of a given category? Do these categorical differences produce differences in discrimination?

To answer these questions, we must first ask what we know about how humans perceive shapes. Do shapes have properties that suggest something about how they're represented? A study by Op de Beeck, Torfs, and Wagemans (2008) using fMRI imaging found that subjective human judgments of shape similarity were highly correlated with the similarity between neuronal response patterns in the occipitotemporal cortex, even when subjectively similar shapes had very different retinotopic representations, even though both the occipitotemporal cortex and more retinotopic brain regions (V1, V2, V4, etc.) have been implicated in the organization of perceptual objects (Altmann, Bühlhoff, & Kourtzi, 2003). This suggests that human judgments of similarity can tell us something about how the brain stores and represents shape information. Barenholtz and Tarr (2008) noted that many objects in the real world can articulate while retaining their object identity; a hand is still a hand whether the fingers are straight or bent. They theorized that objects were modelled in the brain as sets of rigid parts, and they found

that human shape similarity judgements were consistent with such a model, where shapes that could be seen as different articulations of the same set of rigid parts were viewed as highly similar relative to shapes that could not be seen as having the same rigid parts. This suggests that the brain uses a part-based shape representation. Graf and Bülthoff (2003) found that shapes that transforming a shape by blending it with another shape rendered it a worse member of its category in a monotonic fashion, such that the more transformed a shape was, the less categorical it was judged as being, which suggests that the mental representation has some sort of distance-like property. However, shape similarity is not always a true distance metric. In a true distance metric, the distance from point A to point B must be the same as the distance from point B to point A, but Han, Close, and Graf (2009) demonstrated that similarity violates this requirement. By showing videos of a shape transforming into another shape, and then asking subjects to rate the similarity of the first shape to the second, Han et al found that shape similarity judgements were affected by which shape was shown first. The similarity of A to B was not the same as the similarity from B to A. Gauthier and Tarr (1997) also found that familiar, well-studied shapes were sensitive to part configuration in much the same way faces are, such that well-studied shapes are more difficult to categorize when viewed from unusual configurations or orientations, which suggests that familiar shapes may have special properties distinct from those of unfamiliar shapes.

But how do we take this information and apply it to determine something concrete about the mental representational space? To do that, we turn to a paradigm

widely used in tasks involving the recognition of familiar objects: morph-space discrimination (Hartendorp, Van der Stigchel, & Postma, 2013; Newell & Bulthoff, 2002; Panis, Vangeneugden, & Wagemans, 2008). Morph-spaces are defined by mathematical transformations between a set of two shapes, which allow the extraction of partially morphed intermediate shapes at any point along a continuous range. Here, we examine morph-spaces defined by taking a weighted average of matched sets of the points for each of the two base shapes, with weights inversely proportional to the morph-space distance from the base shape. Each point on one shape is matched to a corresponding point on the next shape, minimizing the mean squared distance between each pair of points while preserving the order of points along the edges of the shapes. This results in a morph-space consisting of observationally-reasonable transitional shapes between the two base shapes. See figure 1 for an example of shapes extracted from morph-spaces.



*Figure 1.* Shapes extracted from morph-spaces, with base shapes at each end and transitional shapes in the middle. Above: a lamp on the left morphs into a human on the right. Below: a turkey on the left morphs into a truck on the right. Extracted from Hartendorp, Van der Stigchel, & Postma (2013).

Morph-spaces have been used in the past to study categorization effects.

Morph-spaces between familiar objects have been used to investigate the categorical structure of mental representations for such objects, and they have been used to

demonstrate several important categorization properties. Hartendorp et al (2010) tested morph-spaces constructed between two objects that are representative of known, distinct categories (such as in the morph-spaces illustrated in figure 1), and found that most shapes are recognized as members of one of the two categories. However, shapes near the center of the distribution prove more difficult to categorize with certainty, and may not even be recognized as belonging to either category.

Newell and Bulthoff (2002) also tested morph spaces between familiar objects, with an eye towards how easily pairs of objects are discriminated from one another. By measuring difference thresholds, Newell and Bulthoff found that object pairs are more easily discriminated when the two objects are separated by a category boundary (see also Goldstone, 1994; Folstein, Palmeri, & Gauthier, 2012). Gillebert et al (2008) conducted a similar experiment using 2-dimensional morph-spaces (with four base shapes, one at each “corner”) and found that discrimination was easiest for objects that were easily categorized.

By using morph-spaces, these researchers have found a strong relationship between discrimination and familiar object categories, but what does that say about shape? Object categories are defined by more than just shape contours. They’re also defined by functional and semantic associations (Murphy & Medin, 1985; Rehder, 2003; Op de Beeck, Torfs, & Wagemans, 2008; Graf & Folstein, Palmeri, & Gauthier, 2014). Previous morph-space research has shown that discrimination is strongest across well-defined category boundaries (Newell & Bulthoff, 2002), and in parts of the shape space that the individual is very familiar with (usually within well-defined categories;

Hartendorp et al, 2010). However, in spite of these findings, the implications for mental shape representations are still unclear.

To address the fundamental issue of shape representation, another approach is needed. We know from the aforementioned research that discrimination along morph-spaces depends on category boundaries, but the category structure of unfamiliar shapes is not immediately apparent. However, because discrimination depends on category boundaries, it is also possible to infer category boundaries from discrimination patterns. The relationship between category structure and discrimination can be investigated in both directions. Just as past research has investigated discrimination as it relates to known category structures, we can now examine unknown category structures by analyzing discrimination patterns.

In order to assess the mind's representational structure for raw shape properties, we constructed morph-spaces between pairs of unfamiliar shapes (see table 1). Participants were then shown pairs of shapes drawn from nearby points in the morph-space, and asked whether that shape pair was the same or different. By varying the morph-space distances between the two shapes, we were able to estimate a difference threshold for each point along the morph-space.

Using these difference thresholds, we set out to measure discrimination along morph-spaces, first in order to determine whether discrimination was uniform along these spaces, and then to discover what patterns of sensitivity to shape differences exist in the morph-spaces. Our main focus was on understanding what these variations, if

they exist, indicate about the underlying category structure of the brain's representations of shape.

## DESIGN OF THE EXPERIMENTS

To find these difference thresholds, we applied the same experimental procedure to many different sets of stimuli across a number of subjects. Each experiment consisted of a shape discrimination task, the results of which were used to estimate difference thresholds for each subject at various points along the morph space.

Subject were told that they would play a game where they attempted to defend a space station. They would see pairs of shapes moving slowly down a computer monitor, and each shape pair would either consist of two identical shapes or two non-identical shapes. Pairs of identical shapes were friendly aliens that needed to be allowed into the space station, while pairs of non-identical shapes were asteroids or hostile aliens that needed to be shot down before they could endanger the space station. Subjects were instructed to press the spacebar to allow a pair in, or the 'z' key to shoot a pair down. Subjects were also instructed to disregard rotation when comparing shapes, and in some experiments, subjects were also told to disregard relative size. Subjects were told to respond to each shape.

Subjects were then shown these shape pairs, one pair at a time. Each pair appeared as described in the instructions, either identical or non-identical. The pairs would move slowly down the screen, until the subject either shot them down (non-identical judgment) or allowed them in (identical judgment), as described in the instructions. In each pair, one shape was rotated 15 degrees in order to make direct pixel-by-pixel comparisons impossible, forcing subjects to rely on mid-level mental representations of the shapes.

Each experiment drew shapes from a morph-space between two base shapes, which differed between experiments. The morph-spaces were generated by matching points along the edges of each shape, and then taking a weighted average of the coordinate locations of each pair of matched points. The resulting shape were computed as follows:

$$S_{final} = \alpha s_1 + (1 - \alpha)s_2$$

where  $\alpha$  is the morph proportion,  $s_1$  is a  $2 \times n$  array containing the coordinate values of the  $n$  points used to define the first base shape, and  $s_2$  is a  $2 \times n$  array containing the corresponding coordinate values for the second base shape.  $s_{final}$  is a  $2 \times n$  array containing the coordinate values of the points defining the morphed shape. At  $\alpha = 0$ , the morph-space produces the first base shape, and at  $\alpha = 1$  the morph-space produces the second base shape. Intermediate values of  $\alpha$  produce a smooth range of intermediate morphs.

Each experimental session was divided into either 9 or 27 blocks. Each block corresponded to one of nine  $\alpha$ -values along the morph-space (with either 1 or 3 repetitions for each  $\alpha$ -value, depending on the experiment), ranging from  $\alpha = .1$  to  $\alpha = .9$  at intervals of .1. In each block, one shape in each shape pair will always be the shape resulting from that block's corresponding  $\alpha$ -value (so in block  $\alpha = .4$ , one shape in each pair will be the shape produced when  $\alpha = .4$ ).

The second shape in each pair was selected using an adaptive method, which determined the next comparison shape based on the subject's previous responses. After each response, the algorithm adjusted an estimated difference threshold in terms of  $\alpha$ -



values. If the estimated threshold fluctuated by less than .005  $\alpha$  units for 10 trials in a row, the program would accept the most recent estimate and move to the next block. As a result, a block did not contain a set number of trials. The algorithm used here was the Psi method described in Kingdom and Prins (2010).

Blocks were presented in rounds of 9 blocks, one for each  $\alpha$ -value under investigation. Within a given round, the order of blocks was randomized. For the 9-block experiment, subjects saw 1 full round, while in the 27-block experiment, subjects saw 3 full rounds. This resulted in multiple threshold estimates for each  $k$ -value for each subject, all of which were averaged together to estimate the subject's actual threshold at that point in the morph-space.

## PARTICIPANTS

The participants consisted of 49 (32 female) individuals recruited using fliers posted in and around Rutgers university's Busch Campus. Participants were paid \$28 for 2 hours of participation in the experiment, and saw 3 full rounds of trials. The group ranged from 18 to 35 years of age.

## EXPERIMENTS

The subjects were divided into 14 distinct shape conditions. In each of these conditions, subjects saw three full rounds of trials in one of eight different morph-spaces. In addition, the shapes used in each condition were subject to one of four different normalization styles, depending on the condition, in order to determine whether subjects relied on simple area or perimeter comparisons when determining whether shape pairs were identical. The four normalizations were as follows:

**No Normalization:** No normalization was performed on the shapes produced by the morph-space. However, this means that if a shape part was present in both base shapes (such as in the one-part/two-part morphs shown in table 1), that shape part would remain the same size throughout the experimental session. As such, No Normalization might actually be considered Normalized Part Sizes.











**Normalized Area:** Shape pairs were normalized such that they always had the same area. However, this means that the size of any individual part may change throughout the session, in order to keep the overall size of the shapes constant.

**Normalized Perimeter:** Shape pairs were normalized such that they always had the same perimeter. The same caveats from Normalized Area apply.

**Random Scale:** Shapes were randomly scaled to be larger or smaller on each trial, and subjects were explicitly instructed to ignore relative size in their comparisons.

The number of subjects run in each condition is shown in table 2.

Table 1.  
*Listing of Morph-Spaces: Example shapes and corresponding labels*

Sample Shapes	Name
	One-Part/Two-Part primary
	One-Part/Two-Part alternative
	Two-Part/Three- Part primary
	Two-Part/Three- Part alternative
	Two-Part/Three- Part alternative 2
	Two-Part/Three- Part alternative 3
	Circle/Ellipse
	Ellipse/Bent Ellipse
	Ellipse/Peanut
	Two-Part/Three- Part regular

## RESULTS

One-Part/Two-Part Morphs: As seen in figure 2, morphs between one-part shapes and two-part shapes appear highly sensitive to differences in normalization. In the no normalization condition, the primary one-part/two-part morph (see table 1) displays strong sensitivity near the one-part shape, dropping off monotonically as the morph progresses towards the two-part shape. In the normalized perimeter condition, however, this trend reverses, and sensitivity is highest near the two-part shape. In the random scale condition, there are no clear trends of increased sensitivity at any point along the morph-space. However, for the alternative one-part/two-part morph, there was elevated sensitivity near the one-part shape in the random scale condition.

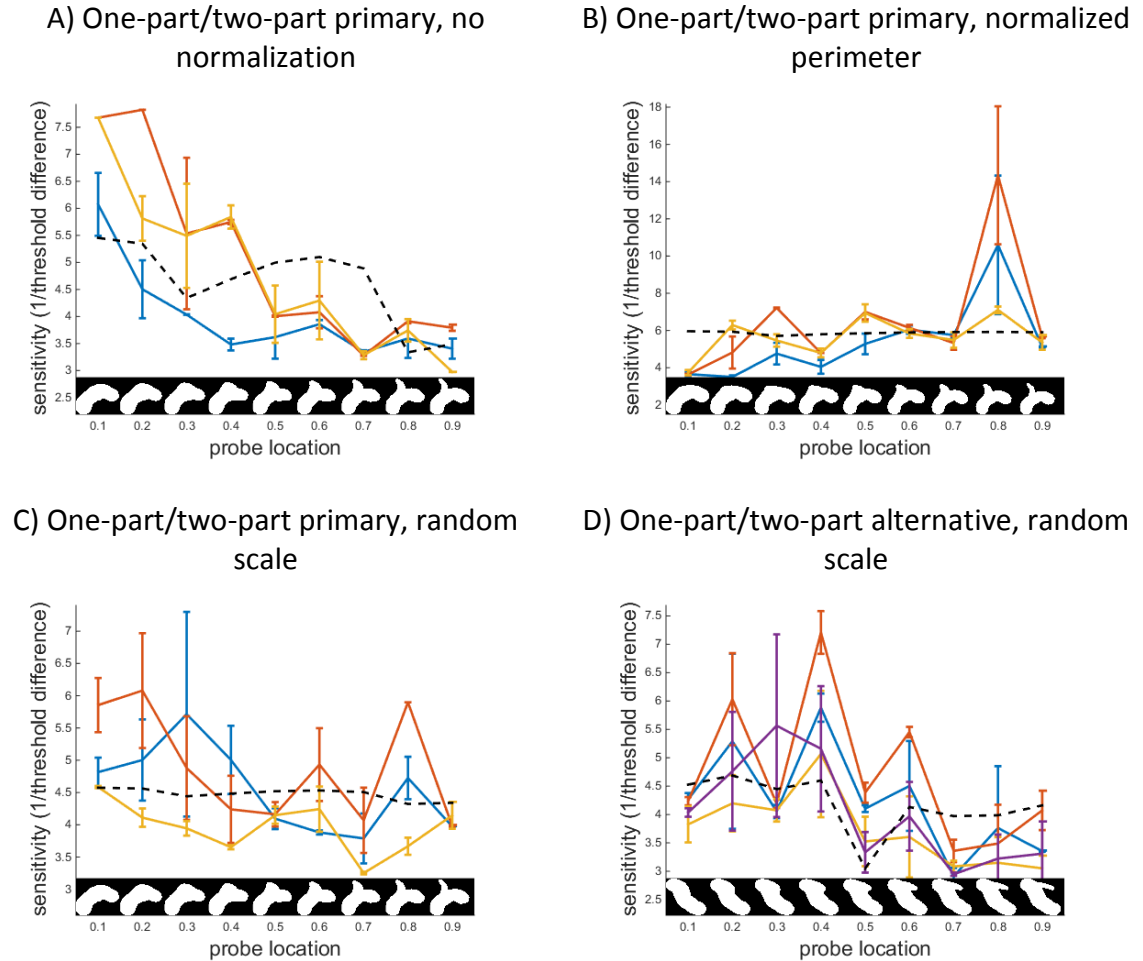
Table 2.

*Experimental Conditions by Morph-Space and Normalization Method*

	No Normalization	Normalized Area	Normalized Perimeter	Random Scale
<b>One-Part/Two-Part</b>	3	--	3	3
<b>One-Part/Two-Part alternative</b>	--	--	--	4
<b>Two-Part/Three-Part</b>	--	2	2	--
<b>Two-Part/Three-Part alternative</b>	--	1	2	--
<b>Two-Part/Three-Part alternative 2</b>	--	--	--	3
<b>Two-Part/Three-Part alternative 3</b>	--	--	--	2
<b>Circle/Ellipse</b>	--	6	2	4
<b>Ellipse/Bent Ellipse</b>	--	3	2	--
<b>Ellipse/Peanut</b>	--	3	--	--
<b>Two-Part/Three-Part regular</b>	4	--	--	--

Two-Part/Three-Part Morphs: As seen in figures 3 and 4, morphs between two-part shapes and three-part shapes appear largely resilient to different normalizations. The primary one-part/two-part morph shows a slight tendency towards lower sensitivity in the vicinity of the two-part shape, and then no overall trend between the middle of the morph and the three-part shape. In the normalized perimeter condition, there is also a suggestion of a bimodal sensitivity distribution, which is seen again in the two-part/three-part alternative morph. In that morph, the sensitivity distribution appears bimodal and generally highest towards the middle of the distribution, and lowest towards the end. Normalization does not appear to have any meaningful effect. In the second alternative two-part/three-part morph, the distribution shows an extremely clear trend of high sensitivity towards the middle of the morph, and low sensitivity near the base shapes. Finally, in the morph between regular one-part and two-part shapes (which was not normalized), there was a slight trend towards increased sensitivity in the middle of the distribution.

Ellipse-Based Morphs: As seen in Figure 5, morphs involving ellipses and ellipse-like shapes, rather than skeletal shapes, showed no variation across normalizations. The circle/ellipse morph showed heightened sensitivity near the circle and monotonically decreasing sensitivity as the morph approached the base ellipse. The ellipse/peanut morph showed the same trend. The ellipse/bent ellipse morph, by contrast, showed heightened sensitivity towards the middle of the morph-space and reduced sensitivity near each base shape.

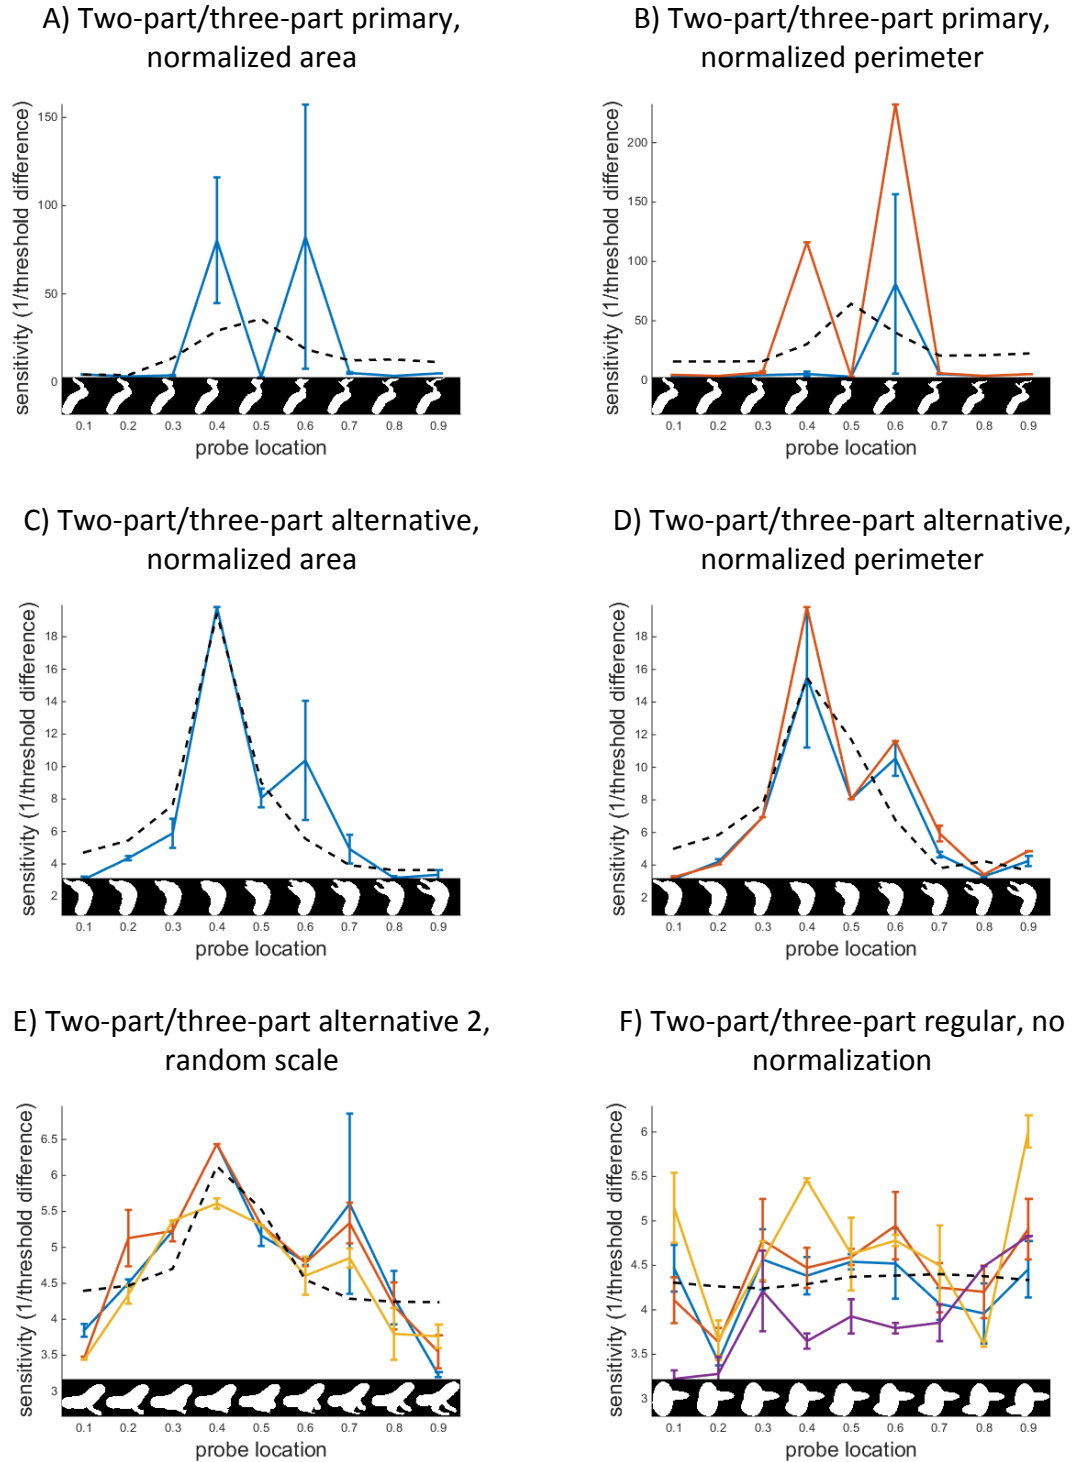


*Figure 2.* Sensitivity (1/difference threshold) curves for each one-part/two-part condition. Each colored line represents a single subject, while the dotted black line represents the predictions of the fitted model, as described in the discussion section. Shapes are shown across the bottom of each graph, morphed in approximate proportion to the probe location. Thresholds for these graphs were computed using normalized distance between matched pixels. Morph distance and pixel overlap were also used to compute thresholds, but did not produce substantially different results.

## DISCUSSION

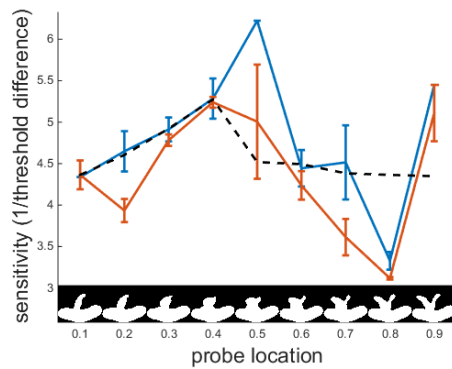
Theoretically, it should be possible to represent discrimination as a simple function of the difference between two shapes in the brain's underlying representational shape space. If two shapes  $a_1$  and  $a_2$  are separated by a distance  $d$  in the representational space, and two shapes  $b_1$  and  $b_2$  are also separated by the same distance  $d$  in this space, then the discriminability of the pair  $a_1$  and  $a_2$  should be the same as the discriminability of the pair  $b_1$  and  $b_2$ . Sensitivity, should vary according to the distance in the representational shape-space, and if the brain uses simple pixel representations or slightly-more-complex contour-point representations, we would expect sensitivity to be a simple function of more-space or pixel distance. However, we can see from the data that this is not the case. If sensitivity were a linear function of either morph-space distance or pixel distance, we would expect human results to be constant when measured in terms of those distance measures. However, as can be seen in figures 2 through 5, this is not the case. If sensitivity were a more complex function of morph-space distance or pixel distance, we would expect a consistent pattern across all conditions, but again, figures 2 through 5 show that this is not the case.





*Figure 3.* Sensitivity (1/difference threshold) curves for the first six two-part/three-part conditions. Each colored line represents a single subject, while the dotted black line represents the predictions of the fitted model, as described in the discussion section. Shapes are shown across the bottom of each graph, morphed in approximate proportion to the probe location. Thresholds for these graphs were computed using normalized distance between matched pixels. Morph distance and pixel overlap were also used to compute thresholds, but did not produce substantially different results.

Two-part/three-part alternative 3, random  
scale



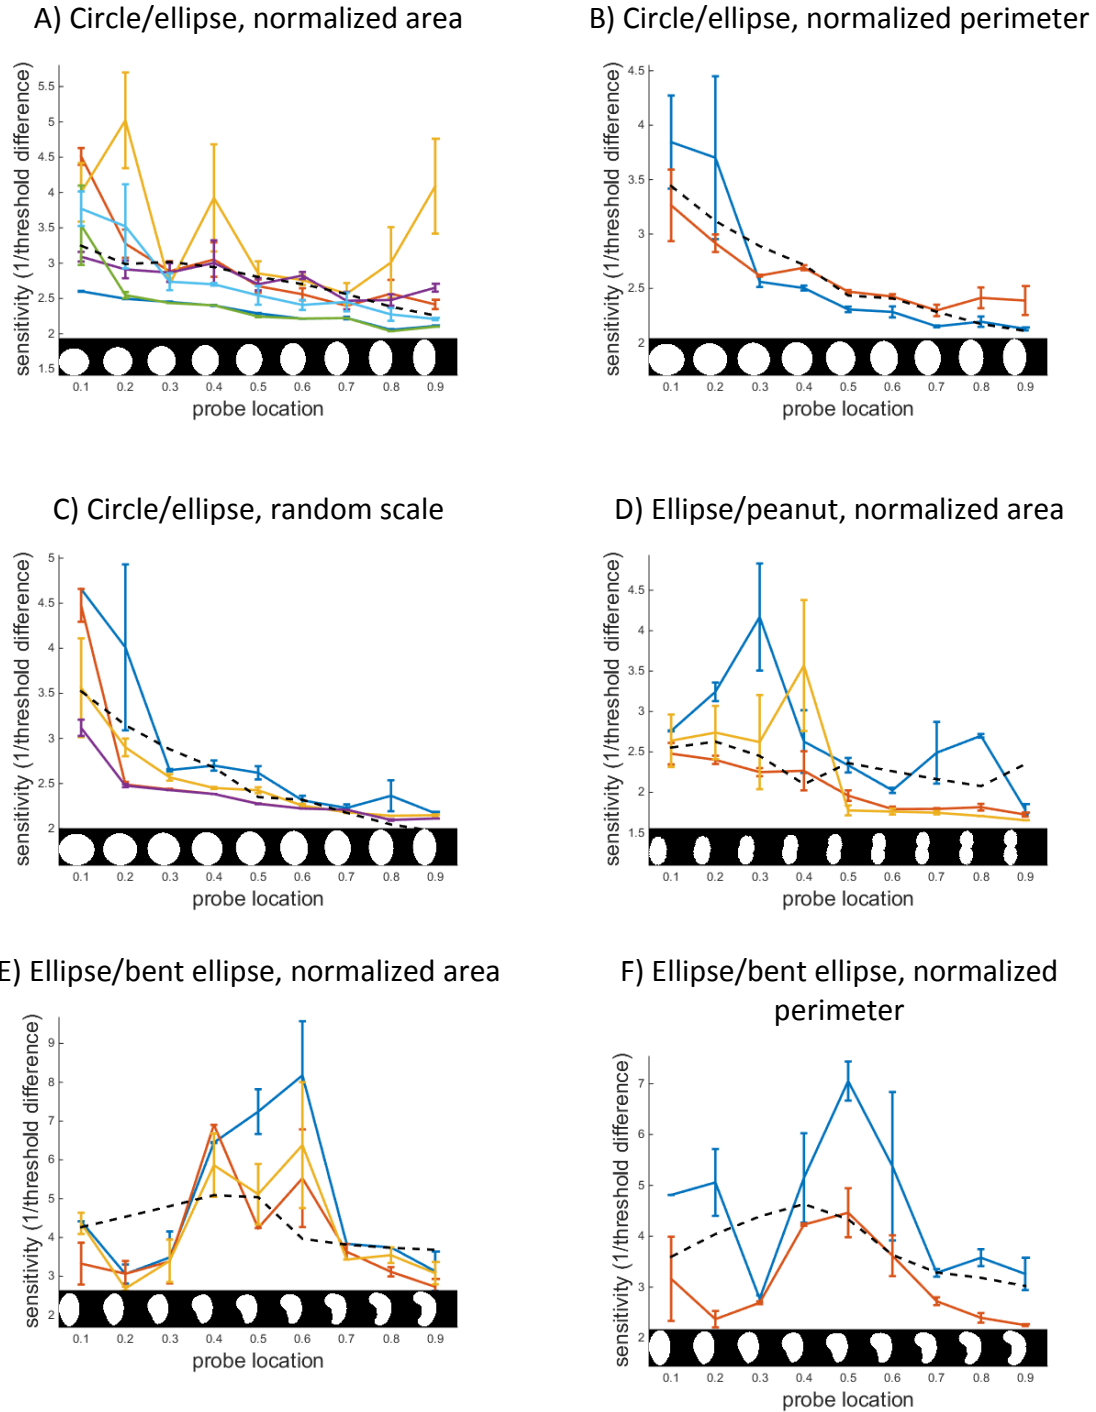
*Figure 4.* Sensitivity (1/difference threshold) curves for the last two-part/three-part condition. Each colored line represents a single subject, while the dotted black line represents the predictions of the fitted model, as described in the discussion section. Shapes are shown across the bottom of each graph, morphed in approximate proportion to the probe location. Thresholds for these graphs were computed using normalized distance between matched pixels. Morph distance and pixel overlap were also used to compute thresholds, but did not produce substantially different results.

The fact that sensitivity is not a simple function of pixel distances suggests that discrimination is based on a more abstracted shape representation, not just on raw pixel distances. Hence, if the brain were simply comparing raw image data to determine whether two shapes were identical, sensitivity would have to be a function of such pixel distances. Since sensitivity is not a function of pixel distances, subjects must be using more abstract shape representations to perform the comparison.

So to explain the results, we turn to models of shape representation itself. Instead of basing sensitivity predictions on raw image properties or mathematically convenient abstractions, we will attempt to model the structure of the representational space. Sensitivity patterns are then predicted by the distance between shapes.


$$d' \propto \text{distance}(A, B)$$

But what is the distance between two shapes? To model this, we will turn to work by Briscoe (2008), in which she found that probabilistic shape skeletons could be



*Figure 5.* Sensitivity (1/difference threshold) curves for each ellipse-based condition. Each colored line represents a single subject, while the dotted black line represents the predictions of the fitted model, as described in the discussion section. Shapes are shown across the bottom of each graph, morphed in approximate proportion to the probe location. Thresholds for these graphs were computed using normalized distance between matched pixels. Morph distance and pixel overlap were also used to compute thresholds, but did not produce substantially different results.

used to predict human similarity judgments between shapes. These skeletons, based on work by Feldman and Singh (2006), consist of sets of axes and rib distributions that probabilistically generate shape outlines. Axes form the core of the shape skeleton, and are probabilistically “inflated” by adding ribs according to distributions over rib length and direction. These rib endpoints then form the outline of the final shape (see figure 6). The shape’s axes, then, typically correspond to the subjective part structure of the final shape.



*Figure 6.* Sample shape skeletons, from Feldman and Singh (2006). *a* depicts the generative process that takes the axis and adds ribs to generate an inflated shape. *b* shows the maximum posterior skeleton for the outline of a human hand. Note that the axes roughly correspond to the subjective part structure of the hand.

Because this generative process is probabilistic, each skeleton has a probability distribution over possible shapes that it can generate. Thanks to Bayes’ Rule, this means that each shape also has a probability distribution over possible skeletons that could have generated it. These two probabilities are related by Bayes’ rule.

$$p(skel|shape) \propto p(shape|skel)p(skel)$$

This probabilistic process allows us to take an approach similar to Ashby and Perrin's (1988) General Recognition Theory, in which more general objects are represented as probabilistic distributions, and the representational distance between two objects is taken as a function of the probability of one object being a member of the other object's distribution. In our model, a key component of the shape representation is considered to be a probabilistic generative function (the shape skeleton) that reproduces the shape with a certain degree of error. Following Briscoe (2008), we hypothesize that distance will be proportional to the cross description length between the two shapes.

$$distance(A, B) \propto cross\_dl_{A,B}$$

The cross description length (*cross\_dl*) of shapes *A* and *B* is a measure of how likely each shape is to be generated by the other shape's most probable skeletal generator (according to the method outlined in Feldman and Singh, 2006). That is, the likelihood of shape *A* given skeleton *B*, and the likelihood of shape *B* given skeleton *A*. All skeletons used in this model are the maximum posterior (MAP) skeletons for their respective shapes.

$$likeli_{AB} = p(shape_A|skel_B)$$

$$likeli_{BA} = p(shape_B|skel_A)$$

The cross description length is then computed by taking the average of the two negative log likelihoods, as shown below.

$$cross\_dl_{A,B} = \frac{-\log(likeli_{AB}) - \log(likeli_{BA})}{2}$$

A cross description length is computed for each of the nine tested locations along each morph-space. For each location, shape  $A$  is the shape at the current morph-space location ( $location_A$ ). Shape  $B$  is chosen as the shape at location  $B$ , as determined by the equation below:

$$location_B = location_A + \Delta M$$

Where  $\Delta M$  is a constant morph-space distance. In this model,  $\Delta M$  was always .1 morph units, 10% of the total distance between the two base shapes used to define the morph. Using this method, the cross description lengths provide a locally-defined “distortion” factor from the morph-space  $M$  to the representational space  $R$ , in which sensitivity is predicted directly by distance. Both  $M$  and  $R$  can be viewed as distortions of one another. If  $M$  is viewed as a flat space (a useful simplifying assumption, since  $M$  is known a priori and  $R$  is not), then  $R$  can be generated by a continuous series of instantaneous distortion factors  $\frac{dR}{dM}$ . Any distance between two points in  $R$  can be derived from the distance between the corresponding points in  $M$  and the distortions  $\frac{dR}{dM}$  along the space between those points.

$$R\_dist_{A,B} \propto M\_dist_{A,B} \cdot \frac{dR}{dM}$$

The cross description lengths computed by the model do not provide instantaneous distortion factors, but they do provide a discrete estimate of the distortion factor over a given  $\Delta M$  window. In this way, the distortion factors produced by the model are an approximation of the true instantaneous distortion factors. In the limit, as  $\Delta M$  goes to 0, the approximation will equal the true distortion factor. In our

model,  $\Delta M$  is not very close to 0, and so it provides a very rough approximation to the true distortion.

$$\frac{\Delta R}{\Delta M} \sim \frac{dR}{dM}$$

By providing an estimate of the local distortion from  $M$  to  $R$ , the model allows us to make predictions about sensitivity based on the distortion. Specifically, the more representational distance is traversed by the same stretch of morph-space, the higher sensitivity should be in that region. After all, the “true” representational distance between two shapes in that region is greater than predicted by the morph-space alone, and representational distance is proportional to sensitivity.

As a note, while the distortion factor is computed from  $M$  to  $R$ , it is trivial to convert from  $M$  to a pixel-distance space, or some other geometrically-defined shape space, because  $M$  and pixel-distance are related by simple geometric calculations that do not involve psychological representations. Thus, by determining the distortion from  $M$  to  $R$ , we determine the distortion from all well-defined geometric shape spaces to  $R$ .

In order to make sense of these distortions, we compute a set of nine cross description lengths (the distortion factors) for each morph-space, one for each of the nine tested locations. Each set of nine cross description lengths is then fitted to the observed human sensitivity values using a simple 2-parameter linear fit, in order to properly position and scale the model predictions to the data. This fitting is necessary because description length does not use the same units as sensitivity. The model posits that cross description length is proportional to sensitivity, and in order to turn that proportionality into a prediction, the linear fit is applied.

This entire process is then repeated for each morph condition. Cross description lengths are computed for each tested location in each morph space, and then linearly fitted to the data to produce the model predictions. Predictions are shown as dotted black lines in figures 2-5. As can be seen in the figures, the predictions give a reasonable approximation to the human data, at least in most cases. In order to affirm this intuition, an AIC was computed for the skeleton-based cross-likelihood model, and for a model using only the pixel distance between the two shapes. See table 3 for the results

Table 3.

*AIC Results by Shape and Normalization Condition*

<b>Shape Condition</b>	<b>Skeletal Model AIC</b>	<b>Constant Model AIC</b>
<b>One-Part/Two-Part</b>		
No normalization	184.3472	197.1168
Normalized perimeter	392.2722 <sup>a</sup>	390.3364 <sup>a</sup>
Random scale	233.1251 <sup>a</sup>	231.7254 <sup>a</sup>
<b>One-Part/Two-Part alternative</b>		
Random scale	321.81	333.8235
<b>Two-Part/Three-Part</b>		
Normalized area	291.0397 <sup>a</sup>	290.5649
Normalized perimeter	602.0987	603.2586
<b>Two-Part/Three-Part alternative</b>		
Normalized area	133.4905	169.9436
Normalized perimeter	270.2524	323.2172
<b>Two-Part/Three-Part alternative 2</b>		
Random scale	130.4771	135.8006
<b>Two-Part/Three-Part alternative 3</b>		
Random scale	182.5402	224.903
<b>Two-Part/Three-Part regular</b>		
No normalization	231.6787 <sup>a</sup>	230.3486 <sup>a</sup>
<b>Circle/Ellipse</b>		
Normalized area	322.0451	352.2819
Normalized perimeter	56.44244	96.79211
Random scale	143.446	225.3872
<b>Ellipse/Peanut</b>		
Normalized area	176.7095	180.053
<b>Ellipse/Bent Ellipse</b>		
Normalized area	240.4451	245.1462
Normalized perimeter	190.0998	196.2669

<sup>a</sup>Skeletal model outperformed by constant model in these conditions



of those comparisons.

As can be seen from the AIC values in table 3, our model outperforms the mathematically simple comparison model in most cases, and performs about as well as the comparison model in the remaining cases (marked in the table). While the fit to the data is far from perfect, our model provides the best prediction of human shape discrimination performance currently available. This suggests that shape skeletons, or a similar generative model, are a crucial part of how the brain represents shape.

## CONCLUSION

We began this research by asking how the internal representational shape space was structured. We probed that representational space by testing human discrimination patterns, which proved to be highly non-uniform across different morph-spaces, but highly uniform across subjects. This implied that humans share a similar representational space for unfamiliar shapes. We then attempted to model this space using a skeletal generative model. This model posited that the representational distance between two shapes was a function of the probabilistic generative process for each shape. Using this model, with minimal fitting, we were able to predict human data better than a simple geometric model. While the model's predictions are not perfect, they are a generally good predictor of unfamiliar shape discrimination.

However, the model is not without its weaknesses. In particular, the model assumes that similarity is symmetric. That is, that the similarity from  $A$  to  $B$  is the same as the similarity from  $B$  to  $A$ . However, this is not generally true of similarity judgments (Tversky & Gati, 1982). The discrimination experiment was not constructed to assess the directional nature of similarity, and so the model could not be made to reflect that directionality for these data. However, it is possible that the asymmetry of similarity is responsible for some of the places where the model fails to accurately predict human sensitivity. While the model could be made to reflect these asymmetries by weighting the two cross description lengths differently, the research necessary to improve the model in that way has not yet been performed.

In spite of this weakness, the model is still highly predictive of human discrimination patterns, suggesting that the underlying representational shape space is well-characterized by some variation of a probabilistic generative model over possible shapes.

WORKS CITED

Altmann, C. F., Bühlhoff H. H., & Kourtzi, Z. (2003). Perceptual organization of local elements into global shapes in the human visual cortex. *Current Biology*, 13(4), 342-349.

Ashby, F. G., & Alfonso-Reese, L. A. (1995). Categorization as probability density estimation. *Journal of mathematical psychology*, 39(2), 216-233.

Ashby, F. G., & Maddox, W. T. (1990). Integrating information from separable psychological dimensions. *Journal of Experimental Psychology: Human Perception and Performance*, 16(3), 598-612.

Ashby, F. G., & Perrin, N. A. (1988). Toward a unified theory of similarity and recognition. *Psychological Review*, 95(1), 124-150.

Barenholtz, E., & Tarr, M. J. (2008). Visual judgment of similarity across shape transformations: Evidence for a compositional model of articulated objects. *Acta Psychologica*, 128, 331-338.

Briscoe, E. J. (2008). *Shape skeletons and shape similarity*. (Unpublished doctoral dissertation). Rutgers University, New Brunswick, New Jersey.

Chellappa, R. (2016). The changing fortunes of pattern recognition and computer vision. *Image And Vision Computing*, doi:10.1016/j.imavis.2016.04.005

Feldman, J., & Singh, M. (2006). Bayesian estimation of the shape skeleton. *Proceedings of the National Academy of Sciences*, 103(47), 18014-18019.

Folstein, J. R., Palmeri, T. J., & Gauthier, I. (2012). Category learning increases discriminability of relevant object dimensions in visual cortex. *Cerebral Cortex*, 23, 814-823.

Folstein, J. R., Palmeri, T. J., & Gauthier, I. (2014). Perceptual advantage for category-relevant perceptual dimensions: the case of shape and motion. *Frontiers in Psychology*, 5, 1394.

Gauthier, I., & Tarr, M. J. (1996). Becoming a “Greeble” expert: Exploring mechanisms for face recognition. *Vision Research*, 37(12) 1673-1682.

Gillebert, C. R., Op de Beeck, H. P., & Wagemans, J. (2008). The influence of categorisation on the perceived shape similarity of everyday objects. *Psychologica Belgica*, 48(4), 261-282.

Goldstone, R. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, 123(2), 178-200.

Graf, M., Bühlhoff, H. H. (2003). Object shape in basic level categorisation. *European Cognitive Science Conference, Osnabrück, Germany, 2003*.

Hahn, U., Close, J., & Graf, M. (2009). Transformation direction influences shape-similarity judgments. *Psychological Science*, 20(4), 447-454.

Hartendorp, M. O., Van der Stigchel, S., & Postma, A. (2013). To what extent do we process the nondominant object in a morphed figure? Evidence from a picture-word interference task. *Journal of Cognitive Psychology*, 25(7), 843-860.

Hartendorp, M. O., Van der Stigchel, S., Burnett, H. G., Jellema, T., Eilers, P. H. C., & Postma, A. (2010). Categorical perception of morphed objects using a free-naming experiment. *Visual Cognition*, 18(9), 1320-1347.

Kingdom, F. A. A., & Prins, N. (2009) *Psychophysics: A practical introduction*. Academic Press. London, UK: Elsevier.

Murphy, G. L., Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92(3), 286-316.

Newell, F. N., & Bulthoff, H. H. (2002). Categorical perception of familiar objects. *Cognition*, 85(2), 113-143.

Op de Beeck, H. P., Torfs, K., & Wagemans, J. (2008). Perceived shape similarity among unfamiliar objects and the organization of the human object vision pathway. *The Journal of Neuroscience*, 28(40), 10111-10123.

Panis, S., Vangeneugden, J., & Wagemans, J. (2008). Similarity, typicality, and category-level matching of outlines of everyday objects. *Perception*, 37, 1822-1849.

Rehder, B. (2003). A causal-model theory of conceptual representation and categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(6), 1141-1159.

Rosch, E. (1973). Natural categories. *Cognitive Psychology*, 4(3), 328-350.

Shepard, R. N., Hovland, C. I., & Jenkins, H. M. (1961). Learning and memorization of categories. *Psychological Monographs: General and Applied*, 75(13), 1-42.

Tversky, A., & Gati, Itamar. (1982). Similarity, separability, and the triangle inequality. *Psychological Review*, 89(2), 123-154.