Vision Research 158 (2019) 189–199

Contents lists available at ScienceDirect

Vision Research

journal homepage: www.elsevier.com/locate/visres

Shape discrimination along morph-spaces

Nathan Destler*, Manish Singh, Jacob Feldman

Department of Psychology, Center for Cognitive Science, Rutgers University, New Brunswick, NJ, United States

ARTICLE INFO

Number of reviewers: 2

Keywords:
Shape
Bayesian
Model
Skeleton
Similarity
Discrimination
Categorization
Medial axis

ABSTRACT

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 category boundaries, reflecting a kind of "categorical perception" in shape space. Here, we use this technique to explore the underlying representation of unfamiliar shapes. 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 potentially important distinctions. Observed discrimination thresholds showed substantial and systematic deviations from uniformity at different points along each shape continuum, meaning that subjects were consistently better at discriminating at certain points along each morph-space than at others. We introduce a shape similarity measure, based on Bayesian skeletal shape representations, which gives a good account of the observed variations in shape sensitivity.

1. Introduction

Shape is an essential aspect of the visual description of objects (Marr, 1982), important for recognition (Biederman, 1987), categorization (Rosch, 1973) and word learning (Landau, Smith, & Jones, 1988). But what object properties actually constitute "shape?" To geometers, shape includes any property of an object's contour that is invariant to position, orientation and scale (Kendall, 1989). But this definition is unhelpful to the psychologist asking what specific features, of the infinite number that in principle could be defined, actually matter to the human visual system. Classical proposals include such simple parameters as aspect ratio (see Feldman & Richards (1998)) and contour curvature (Attneave, 1954; Richards, Dawson, & Whittington, 1988), as well as more complex constructs like parts (Biederman, 1987; Hoffman & Richards, 1984), radial fourier descriptors (Cortese & Dyre, 1996), medial axes (Blum, 1973; Kimia, 2003), scale-spaces (Witkin, 1983), shock-graphs (Siddiqi, Shokoufandeh, Dickinson, & Zucker, 1998), and even more complex descriptors inspired by neural tuning curves in visual cortex (Op de Beeck, Torfs, & Wagemans, 2008; Yamane, Carlson, Bowman, Wang, & Connor, 2008).

In this study we seek to better understand human shape representation by investigating *shape discrimination* along the dimensions of mental shape space. It is well known that certain "qualitative" perceptual boundaries show elevated perceptual discrimination, meaning that human observers are especially sensitive to changes that cross these boundaries and less sensitive to changes that do not, a phenomenon known as *categorical perception* (Harnad, 1987; Goldstone & Hendrickson, 2010). Most famously, linguistic phonemes, though defined by physically continuous phonological parameters, exhibit relatively abrupt categorical boundaries at which native speakers (but not others) have elevated sensitivity to small differences (Harnad, 1987). Categorical perception can occur over other perceptual dimensions as well (e.g. brightness, color, or size), and is usually understood to arise from changes to underlying perceptual discrimination rather than to decision criteria (Notman, Sowden, & Ozgen, 2005). Remarkably, even a single session of category learning on completely novel perceptual categories can induce measurable increases in sensitivity at learned boundaries (Goldstone, 1994).

In the realm of shape, a number of studies have demonstrated sensitization of diagnostic dimensions separating learned shape classes (Folstein, Palmeri, & Gauthier, 2014; Newell & Bulthoff, 2002), and have even identified neural correlates of the process of dimensional

* Corresponding author.
E-mail address: nathan.destler@psych.rutgers.edu (N. Destler).

https://doi.org/10.1016/j.visres.2019.03.002
Received 14 September 2018; Received in revised form 27 February 2019; Accepted 8 March 2019
Available online 22 March 2019
0042-6989/

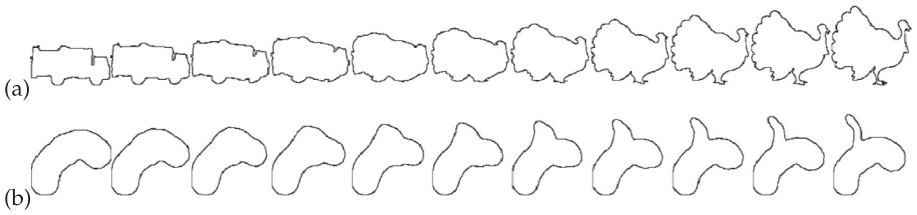

Fig. 1. Examples of shape morph-spaces. Each shape is a linear mixture of the shapes at the poles. (a) A morph between two distinct nameable categories (truck/turkey) similar to those used in previous studies. (b) A morph between two "unfamiliar" shapes similar to those used in our experiments.

modulation (Folstein, Palmeri, Van Gulick, & Gauthier, 2015; Gillebert, Op de Beeck, Panis, & Wagemans, 2009; Sigala & Logothetis, 2002). A common technique in these studies to create a set of shapes connecting one category and another by *morphing* an example from one category (e.g. *truck*) with one from another (*turkey*, Fig. 1). The shapes in a morph space are constructed using weighted combinations of vertices drawn from the two base shapes, with each vertex in the morphed shape consisting of α times a vertex from the first shape and $1 - \alpha$ times a corresponding vertex from the second. This yields a progression of shapes that vary in appearance from one to the other, modulated by the mixing weight α. Shape morph spaces have a variety of uses (see for example Hartendorp, Van der Stigchel, Wagemans, Klugkist, & Postma (2012), Panis, Vangeneugden, & Wagemans (2008) and Wilder, Feldman, & Singh (2011)), and are particularly useful for constructing novel dimensions that modulate category membership.

More specifically, several studies have demonstrated elevated discrimination near the boundaries separating distinct recognizable shape categories. For example, Newell and Bulthoff (2002) constructed morph spaces between two familiar nameable object categories (e.g. truck and turkey), and found that shape discrimination in an ABX task was elevated near the midpoint $\alpha = .5$ of the space, i.e. near the classification boundary between the two shape classes. That is, subjects were better able to detect small differences that crossed the category boundary than those that did not, demonstrating a kind of categorical perception for shape.

In our study, we used a similar technique to study the underlying representation of shape itself, separate from the problem of shape recognition. Previous studies have generally been limited to overtly categorical decisions, either involving previously familiar categories (e.g. truck vs. turkey) or artificial categories induced by laboratory training. But as described above, the underlying mental parameterization of shape space is poorly understood, and we wondered whether patterns of shape discriminability might help reveal some of its properties. In effect, our goal was to discover the *implicit* categorical structure inherent in the representational space itself, by using variations in discriminability to uncover the "seams" in the mental fabric of shape.

Hence in our study, we created morph spaces between "meaningless" shapes (shapes not readily placed into familiar categories) chosen so as to reveal underlying shape dimensions. Our shapes morphs involved artificially constructed comparisons between "minimal pairs" of shapes—shapes in which only one narrowly focused type of shape change separates one base shape from the other, with the morph space smoothly varying between them. Shape changes like "truck" vs "turkey" involve numerous shape differences, consistent with the focus in those studies on previously established named categories. In contrast, our goal was to choose base shapes illustrative of basic shape features, in order to understand the implicit categorical structure (if any) inherent in the shape representation space. The result is a set of morph spaces each of which exemplifies one basic shape transformation. We then tested shape discrimination at points along each morph spectrum, without any overt categorization training or classification task. Though the technique can be applied in principle to any pair of shapes, we tried to choose shape pairs of particular psychological interest, like symmetrical vs. asymmetrical and single-part vs. two-part shapes. The question is then whether (and how) discriminability varies along the morph space—that is, whether differences of a fixed magnitude along a given morph continuum show *unequal* discriminability. For example, certain characteristic shape transitions might show categorical-perception-like elevations in sensitivity to small differences, or other systematic patterns.

Broadly speaking, our goal was to find variations (non-uniformities) in discriminability at points in morph-spaces, that is, points where small shapes differences are detectable to a degree that is disproportionate to the "objective" magnitude of the shape change. Of course, there is no uniquely neutral way to measure the objective degree of shape change. Perhaps the most plausibly neutral way to quantify physical shape change is the total change in vertex positions (after suitable normalization), which we will refer to as the *vertex displacement*, which quantifies how much the geometry of the outer boundary has moved from one shape to another without regard to the nature of the change. The central question in our study was how detectable different *kinds* of shape change—the introduction of a new part, the breaking of a shape symmetry, and so forth—actually are, relative to the magnitude of pixel change actually entailed. We anticipated, as suggested by the results of Denisova, Feldman, Su, and Singh (2016) using a very different methodology, that the nature of the shape change makes a substantial difference to the perceived magnitude of difference. Hence in what follows we are careful always to express detectability in terms of this neutral measure of physical shape change, so that any differences we do find cannot be explained in terms of simple changes in contour vertices, nor as artifacts of the morphing procedure. We are similarly careful to test a variety of different shape normalization procedures, holding area or perimeter constant, randomly scaling the objects, or leaving them unnormalized. We tested these different normalizations in part because the literature is not clear on which aspects of shape scale might affect similarity, and in part to ensure that our results were not artifacts of a particular scale effect. Note that this sort of careful control of shape change is not typical of previous studies that focused primarily on boundaries between distinct nameable shape categories.

More fundamentally, we postulate that subjects' ability to distinguish two shapes generally reflects the degree to which their underlying representations are different—that is, that shapes are discriminable to a degree that corresponds to their dissimilarity in an appropriate shape similarity metric. The main question, then, is the nature of the metric (Edelman, 1998). For example, we might expect "non-accidental" featural differences (such as parallelism or collinearity, Binford (1981), Feldman (2007), Lowe (1987), Wagemans (1992)) to be treated relatively categorically (Feldman, 1997), leading to jumps in dissimilarity as these features are modified, with concomitant increases in discriminability out of proportion to the quantity of physical shape change (Amir, Biederman, & Hayworth, 2012; Kayaert, Biederman, & Vogels, 2003; Kubilius, Wagemans, & de Beeck, 2014). However, non-uniformities in the perceptual shape space need not correspond to intuitively non-accidental properties. For example, the difference between 1- and 2-part shapes might be perceived relatively qualitatively, and thus correspond to local peak in discriminability, but the number of parts has not generally be regarded as a non-accidental feature in existing literature. Our goal is to identify a measure of shape dissimilarity that gives a good account of the observed variations in human shape discrimination *regardless* of their source, whether from non-accidental differences or some other aspect of shape representation.

2. Experiments

In what follows, we applied the same method to a series of shape morphs, each of which was chosen in order to illuminate some aspect of shape representation. In each experiment, subjects were asked to discriminate shapes at various points along the morph space, using an adaptive procedure in order to estimate the difference threshold required to discriminate shapes at each point in the morph space. The method was identical in each of the following experiments except for the choice of base shapes S_1 and S_2.

2.1. General task

Participants were told that they would play a game where they attempted to defend a space station. On each trial, they would see two shapes moving slowly down a computer monitor, which were either identical or non-identical. Identical shape pairs were friendly aliens that needed to be allowed into the space station, while non-identical pairs 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 instructed to disregard orientation (rotation) when comparing shapes, and in some experiments, subjects were also told to disregard relative size (see discussion of normalization, below). In each pair, one shape was rotated 15 degrees relative to the other in order discourage a strategy based on simple pixel-by-pixel comparisons, forcing subjects to rely on intermediate-level mental representations of the shapes and ensuring that any influence of mental rotation was held constant.

Each experiment used shapes drawn from a different morph space defined by the choice of base shapes S_1 and S_2. Morph spaces were generated using the following procedure. Each base shape was a polygonal approximation to a smooth shape consisting of N (= 1000) 2D points, represented as an $N \times 2$ matrix. Base shapes were normalized with respect to area and aligned in terms of the orientations of their common parts (where applicable) so the morphing procedure generally did not introduce any variations in either area or orientation. Each morphed shape at position α along the morph spectrum consisted of the polygon $\alpha S_1 + (1 - \alpha)S_2$, with values $\alpha = 0$ and $\alpha = 1$ yielding respectively base shapes S_1 and S_2, and intermediate values yielding a smooth continuum of intermediate shapes. Morph spaces used in the experiments are visible on the abscissas of the plots shown below.

Each experimental session was divided into either 18 or 27 blocks. Early subjects saw 18 blocks, as initial pilot experiments had suggested that this was the most that could be run in the hour allotted. But after the first five subjects it became apparent that more data could be collected without exceeding the time limit, so subsequent sessions were expanded to 27 blocks. As no feedback was given as to subjects' performance, and no learning effects appeared in the data, subjects from both groups were incorporated into the final data analysis.

Regardless of the total number of blocks seen by the subject, each block corresponded to one of nine α-values along the morph space (with either 2 or 3 repetitions for each α-value, depending on the subject, as above), ranging from $\alpha = .1$ to $\alpha = .9$ at intervals of.1. In each block, one shape in each shape pair was always at the given α value, and the second shape in each pair was selected using the Psi method described in Kingdom and Prins (2010) until a 75% discrimination criterion was met. If the estimated threshold fluctuated by less than .005 units for 10 trials in a row, the program would accept the most recent estimate and move to the next block. The result was an estimate of a shape discrimination threshold at each α value, which we plot as sensitivity (1/threshold expressed in units of 1/pixels).

The order of blocks (α values) was randomized. Subjects saw one full round every 9 blocks: two rounds for subjects in the 18-block group, and three for each in the 27-block group. This procedure resulted in multiple threshold estimates for each α value for each subject, which were averaged to yield mean thresholds at each point in the morph space.

2.2. Participants

The participants consisted of 49 individuals (32 female), ranging from 18 to 35 years of age, recruited using fliers posted on the university campus, and paid for their participation in the experiment. Informed consent was obtained from all participants, and all research described here was conducted under the supervision of the Rutgers Institutional Review Board and in accordance with the Code of Ethics of the World Medical Association (Declaration of Helsinki).

2.3. Conditions

The subjects were divided into 14 distinct shape conditions, based on the morph space used and the nature of the size normalization, the latter comprising several different methods for preventing subjects from using size variations to distinguish shapes. In some cases, we normalized *area*, so both base shapes and morphs had equal area, and in other cases *perimeter*. (It is not generally possible to normalize both area and perimeter at the same time.) In several cases, one base shape was constructed by adding a part to the other base shape (or similar transformations), which meant that normalizing area or perimeter would have had the effect of shrinking the base part of the shape with the added part; so in some cases we used *no normalization*, so that corresponding parts (e.g. the base part) were equal in size. Finally, another approach involved preventing subjects from using object size as a useful discrimination cue by using *random scale* (each object randomly expanded or shrunk by a small ratio) with subject instructed to ignore object size. The number of subjects included in each condition is shown in Table 1.

Our choices of base shapes and resulting morph spaces, illustrated in Fig. 2, reflected a variety of fundamental issues in shape representation. Part-based shape representations (Biederman, 1987; Hoffman & Richards, 1984) posit a "qualitative" difference between shape representations with distinct numbers of parts. Exactly how shapes are

Table 1
Number of subjects used in each of the morph conditions.

	No normalization	Normalized area	Normalized perimeter	Random scale
1-part/2-part, Set A	3	–	3	3
1-part/2-part, Set B	–	–	–	4
2-part/3-part, Set A	–	2	2	–
2-part/3-part, Set B	–	1	2	–
2-part/3-part, Set C	–	–	–	3
2-part/3-part, Set D	–	–	–	2
circle/ellipse	–	6	2	4
ellipse/bent ellipse	–	3	2	–
ellipse/peanut	–	3	–	–
2-part/3-part regular	4	–	–	–

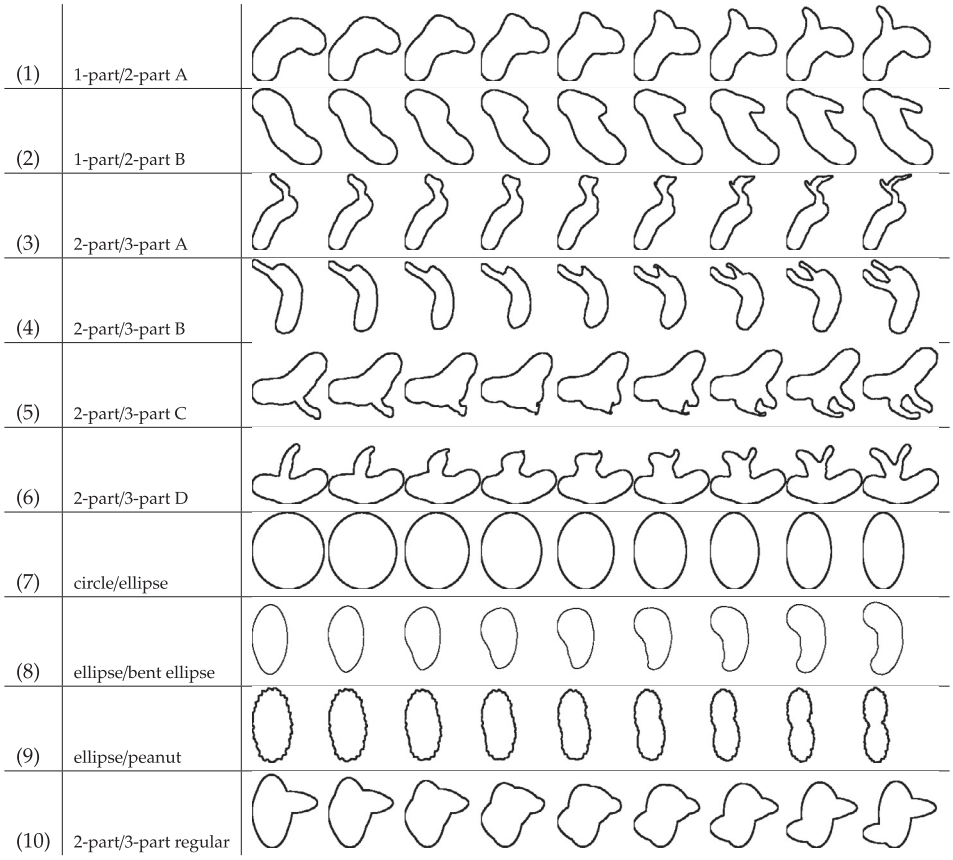

Fig. 2. Morph spaces used in the experiments.

decomposed into parts is the subject of vigorous debate (e.g. de Winter & Wagemans (2006), Feldman & Singh (2006), Singh & Hoffman (2001), Singh, Seyranian, & Hoffman (1999)). But all models assume a psychologically significant distinction between a shape with and without a particular secondary part, such as a torso with and without a particular limb (Siddiqi, Tresness, & Kimia, 1996). Accordingly we included in our morph spaces several spaces in which base shapes differed in part structure, including several versions of 1- vs. 2-part shapes and several of 2- vs 3-part shapes. In each of these conditions we randomly constructed a base shape with the desired number of parts, and then "added" an additional axial part to construct the other base shape. The resulting morph space varies smoothly between the 1-part shape and the 2-part shape (or respectively the 2- and 3-part shape) incorporating a sequence of "intermediate forms" in which the additional part emerges and eventually becomes perceptually distinct. The discrimination experiment then tests sensitivity to small shape changes at various points along this progression, examining whether for example sensitivity gradually increases as the part emerges, or gradually decreases, or peaks at some intermediate threshold of part salience (Hoffman & Singh, 1997).

Similarly, Condition 9 investigates the shape transformation between a simple convex shape (again an ellipse) and a "peanut" featuring concavities on either side. Like the addition of a new limb, this transformation alters the qualitative configuration of parts, albeit in a somewhat different way than investigated in Conditions 1–8. Again the main question is whether the shift to a new part configuration entails a change in sensitivity, perhaps at some perceptual threshold separating the cases, or perhaps monotonically as the transformation progresses.

Another important shape property, tested in Condition 7, was *symmetry*. It is well established that the visual system is very sensitive to shape symmetries, including bilateral symmetry (Palmer & Hemenway, 1978; Wagemans, 1995) as well as various kinds of rotational and transformational symmetries (Chen & Chen, 1987; Wagemans, Van Gool, Swinnen, & Van Horebeek, 1993). Accordingly, we included cases where one base shape exhibits symmetry (e.g. circle) and the other "breaks" the symmetry (e.g. ellipse, in which the circle's rotational symmetry has been broken). The morph space thus exhibits a progression from symmetric to asymmetric (or, more strictly, from symmetric with respect to several transformations to symmetric with respect to one fewer transformation). The discrimination experiment then tests whether this loss of symmetry entails a progressive change in sensitivity to shape change.

Finally, Condition 8 tested the perceptual significance of **bend,** a simple transformation that plays a role in prominent shape theories (Sprote & Fleming, 2016). For example, in Biederman (1987)'s influential theory of recognition-by-components, straight and bent versions of certain shapes are treated as qualitatively distinct primitive shape units. To explore this potentially categorical distinction, the base shapes in Condition 8 are a simple ellipse and a "curved" (banana-shaped) ellipse, with the morph space encompassing a progression of increasing axial curvature. Again the question is whether points along this spectrum entail differences in sensitivity to shape change.

In several cases, in order to avoid drawing conclusions based on what might be accidental features of the chosen shapes, we ran several different (randomly generated) base shape pairs from the same category, designated Set A, Set B and so forth. (Random multipart shapes were generated by randomly choosing axes of the desired part structure and then stochastically generating contours from them using the generative model described in Feldman & Singh (2006).) For example we ran several different size normalizations for the same 1- vs. 2-part morph space (Set A), and then created a different morph space from a different random set of 1- vs. 2-part base shapes (Set B) and ran it as well in order to avoid pinning conclusions on accidental properties of the chosen shapes.

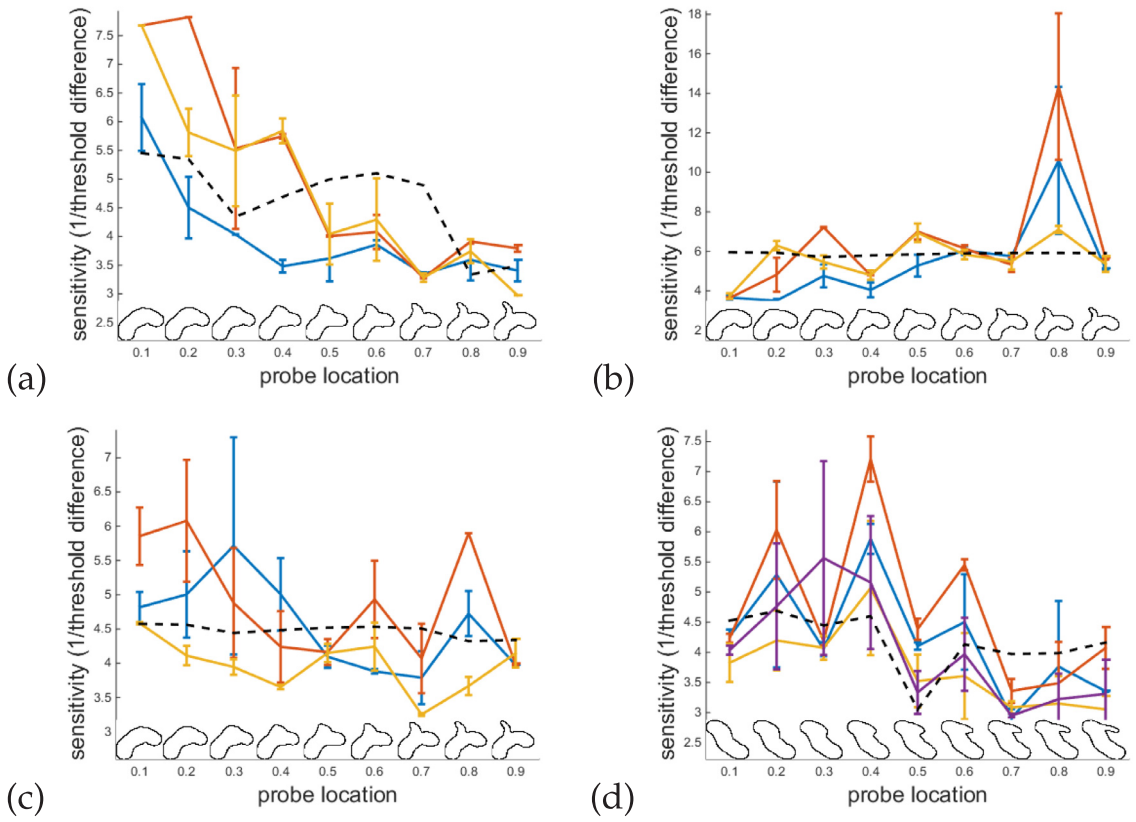

Fig. 3. Results for 1- vs. 2-part shapes, showing sensitivity (1/threshold) as a function of position in morph space, including (a) Set A, no normalization (b) Set A, normalized perimeter (c) Set A, random scale (d) Set B, random scale. Morph spaces are displayed on abscissae. Distinct colors are distinct individual subjects; dotted black line is cross-likelihood similarity model.

2.4. Results

Figs. 3–5 show plots of sensitivity (1/threshold) as a function of position in the morph space, for each morph space and size normalization employed. Each plot shows results for several different individual subjects (colored lines with ±1 s.e. error bars), along with the predictions of a similarity model explained below.

2.4.1. One-part/two-part morphs

As can be seen in Fig. 3, 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 Fig. 2) 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.

2.4.2. Two-part/three-part morphs

As shown in Fig. 4, morphs between two-part shapes and three-part shapes appear largely resilient to different normalizations. The primary two-part/three-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 multimodal 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.

2.4.3. Ellipse-based morphs

Fig. 5 shows results for the remaining morph spaces, each of which pits an ellipse against a simple transformation of an ellipse. These results showed little 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 a similar 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.

3. Discussion

As is plainly visible in Figs. 3–5, subjects showed substantial non-uniformities in shape sensitivity over morph spaces. That is, sensitivity to shape change was higher in some regions of shape space, and lower in others, rather than being constant over the morph mixing parameter α. In most cases, the patterns of variation in sensitivity were qualitatively consistent among subjects; even when subjects differ in absolute sensitivity values, they tend to follow similar patterns with approximately constant offsets. We display individual subject data in order to make clear the degree of agreement among subjects, and in what follows we avoid analyses that depend on aggregating data among subjects, which can give a false impression of a consistent trend.

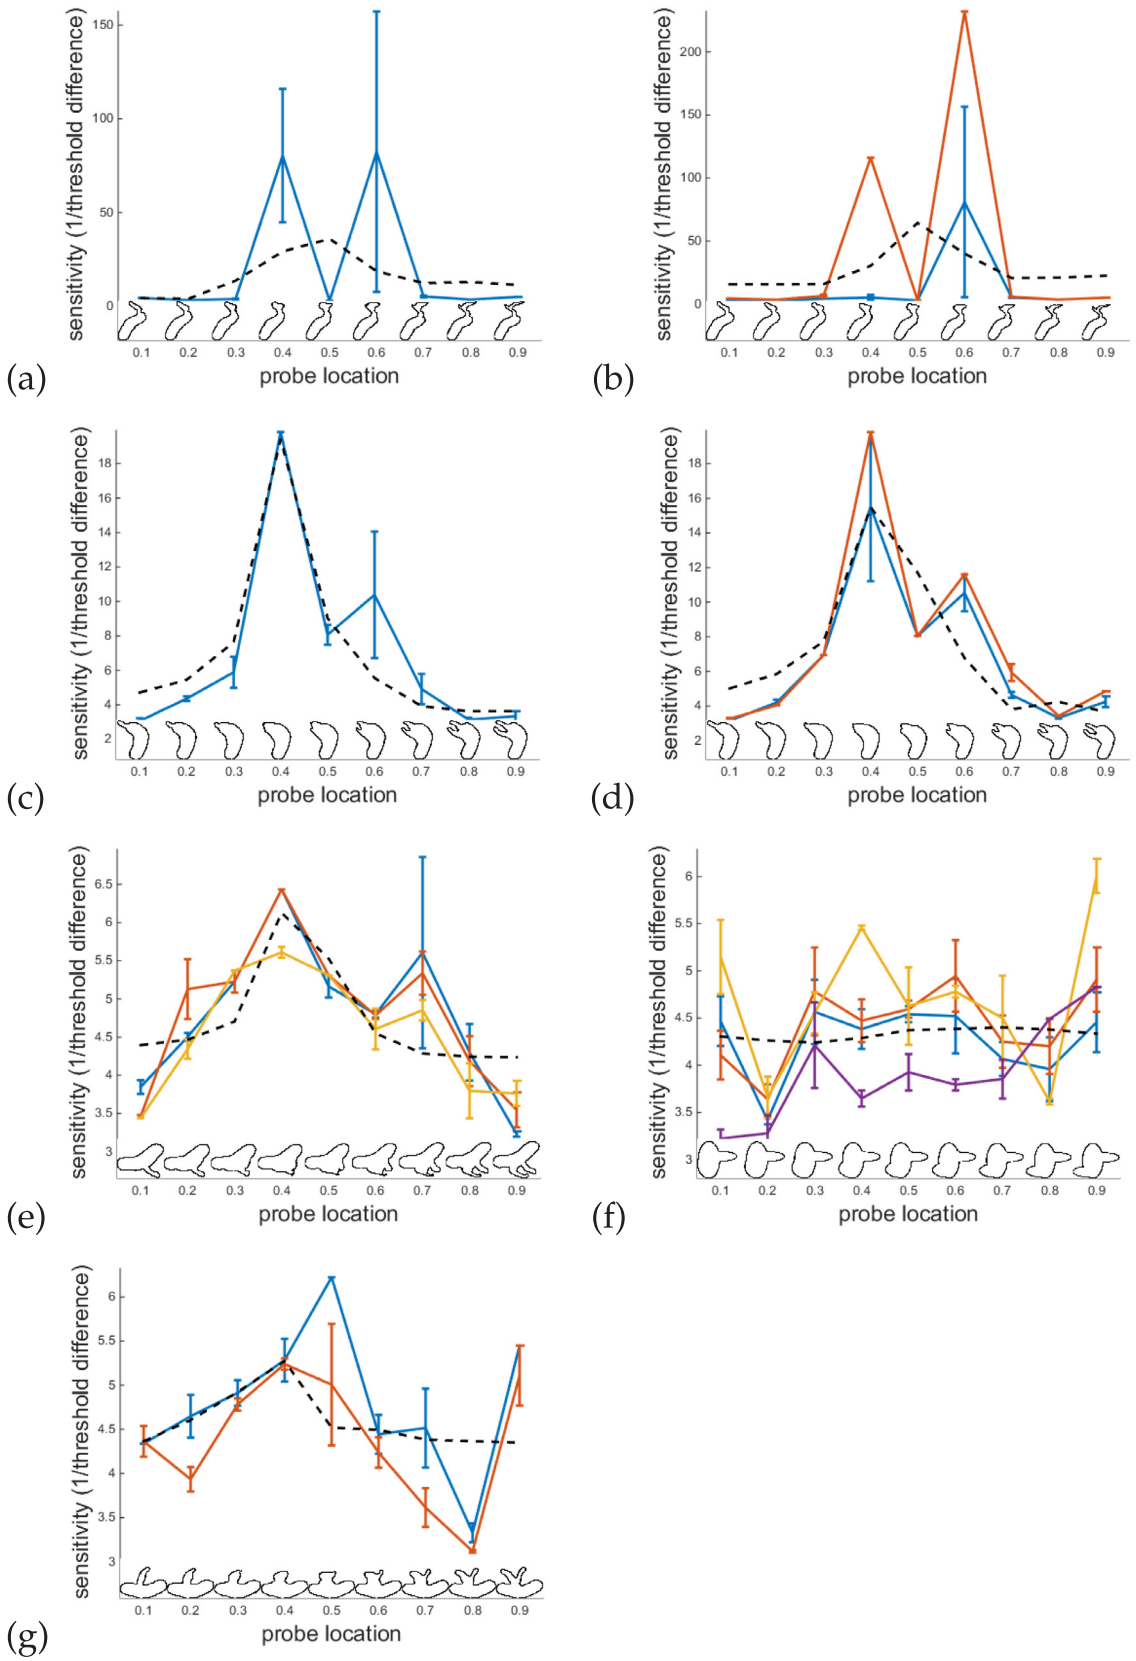

Fig. 4. Results for 2- vs. 3-part shapes, showing sensitivity (1/threshold) as a function of position in morph space. Panels show (a) Set A, normalized area (b) Set B, normalized perimeter (c) Set B, normalized area (d) Set B, normalized perimeter (e) Set C, random scale (f) Set D, no normalization (g) Set C, random scale. Morph spaces are displayed on abscissae. Distinct colors are distinct individual subjects; dotted black line is cross-likelihood similarity model.

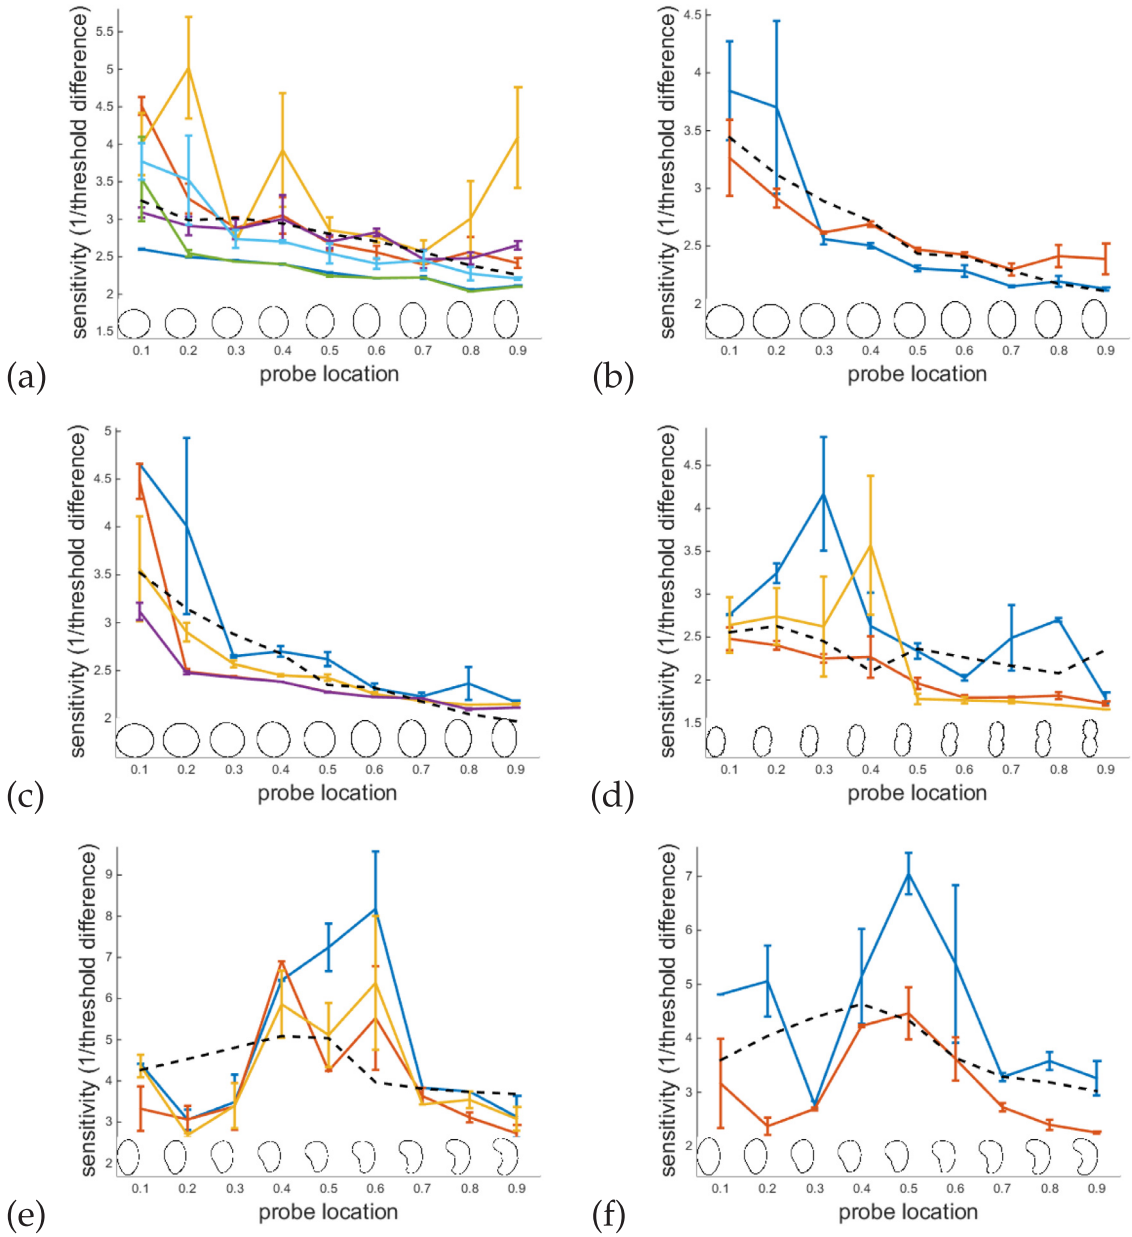

Fig. 5. Results for ellipse-based conditions, showing sensitivity (1/threshold) as a function of position in morph space. Panels show (a) circle/ellipse, normalized area (b) circle/ellipse, normalized perimeter (c) circle/ellipse, random scale (d) ellipse/peanut, normalized area (e) ellipse/bent ellipse, normalized area (f) ellipse/bent ellipse, normalized perimeter. Morph spaces are displayed on abscissae. Distinct colors are distinct individual subjects; dotted black line is cross-likelihood similarity model.

However, while the patterns of variation in sensitivity were often consistent from subject to subject, they vary substantially from morph space to morph space. In some cases subjects showed enhanced sensitivity near a presumed categorical boundary (e.g. 2- vs. 3-part shapes, Fig. 4c,d,e; ellipse vs. bent ellipse, Fig. 5e,f). In other cases subjects showed monotonic increases or decreases in sensitivity along the morph space, culminating in maximal sensitivity at one extreme (e.g. circle vs. ellipse, Fig. 5a-c). In other cases the data were too noisy to detect a clear pattern. In the modeling section below, we aim to provide a unified account of all these patterns using a single consistent shape similarity metric.

Critically, the observed variations in sensitivity *cannot* be explained as artifacts of the morphing procedure. Recall that sensitivity is the reciprocal of the threshold to detect shape differences, expressed in pixels (units of vertex displacement between aligned shape)— the closest thing we have to an "objective" measure of physical shape change. (In other analyses we have also plotted sensitivity expressed as the reciprocal of morph distance $\Delta\alpha$, which yielded very similar patterns of nonuniformity.) That is, the plots show how much "physical" change a shape must undergo before the difference becomes detectable to our subjects, which does *not* in any way depend on the procedure we used to generate the shapes. The purpose of the morphing procedure was simply to map out what *directions* in shape space (i.e., what shape transformations) are particularly detectable and undetectable. The main point is that equal shape differences correspond to very unequal degrees of discrimination depending on where they fall in shape space, and exactly *how* they change the shape at that point.

Several of the specific trends corroborate earlier findings. For example, in the circle/ellipse morph (Figs. 5a–c), discrimination is elevated near the circle and diminishes progressively as the morph

approaches the ellipse, which breaks the circle's rotational symmetry. This bears out the finding by Liu and Kersten (2003) that shape discrimination is better for symmetric than asymmetric shapes. Similarly, discrimination between two- vs. three-part part decompositions appears relatively "categorical" (Fig. 4), showing elevated sensitivity at the boundary separating two and three-part forms. This is consistent with earlier findings suggesting shape changes that involve qualitative changes to part structure are more detectable than those that do not (Barenholtz, Cohen, Feldman, & Singh, 2003; Cohen, Barenholtz, Singh, & Feldman, 2005; Denisova et al., 2016). The pattern for one- vs. two-part shapes is less consistent.

Evidently, the trends visible in the plots implicate a variety of specific shape processes, and further experimentation with other morph classes would doubtless reveal more. However rather than treat each of these patterns as a distinct phenomenon unto itself, our goal is to find an overarching account that explains all the observed variations in shape sensitivity. To this end, below we introduce a single unifying shape similarity measure intended to collectively account for the patterns observed in all the various shape morphs, including both those that conform to previously known shape biases as well as those that are not so readily explained.

3.1. Shape similarity

Our larger goal is not simply to document non-uniformities in shape discrimination, but to explain where they come from. The basic idea is that shape discriminability reflects dissimilarity in the underlying mental shape space—that is, that shapes appear distinct to a degree proportional to the dissimilarity of their mental representations. In order to establish this, we need to introduce a shape similarity measure and show that it effectively predicts our discrimination data.

In the categorization literature, similarity between two items x and y is often assumed to decay exponentially with distance in the underlying shape space,

$$sim(x, y) \propto e^{-cD(x,y)}, \tag{1}$$

where $D(x, y)$ denotes distance in some metric space, and c is a decay parameter modulating how quickly similarity falls off with distance (Shepard, 1957). However, it should be clear that this framework cannot account for our data if D is assumed to be proportional to vertex displacement, since this would predict uniform plots (because Eq. (1) implies that a fixed step size D corresponds to a constant magnitude of dissimilarity, but the observed discrimination is not uniform). Instead, we need a shape similarity measure that reflects the underlying representation of shape, and thus maps different kinds of shape change to different degrees of dissimilarity.

To provide such a measure, we turn to the probabilistic model of shape skeletons introduced in Feldman and Singh (2006). This account of shape representation assumes that shapes are generated by skeletal structures via a process of stochastic lateral growth, and aims to estimate the skeleton *most likely* to have generated the observed skeleton. The result is a skeletal representation of shape, akin to the medial axis representations introduced by Blum (1973) and Blum and Nagel (1978), and later investigated by numerous others (e.g Katz & Pizer (2003), Kimia (2003), Kovács, Fehér, & Julesz (1970), Siddiqi, Shokoufandeh, Dickinson, & Zucker (1999)), except reconceived probabilistically. The probabilistic model assumes a prior over skeletons $p(S)$ (Fig. 6a), which penalizes skeletal complexity, because each additional axial branch, or each additional bend in an axis, entails some probabilistic cost; and a likelihood $p(x|S)$ (Fig. 6b) which defines the probability with which any given skeleton S will generate a given shape x, and thus quantifies the fit between a given shape and skeleton. The MAP (maximum a posteriori) skeleton maximizes the product $p(S)(x|S)$ of the prior and likelihood, jointly optimizing the simplicity of the skeleton and the fit to the shape. Unlike the classical Blum medial axis transform, the MAP skeleton typically corresponds to an intuitive representation of shape with one axis per perceptual part, and has been supported by empirical data in a number of domains (e.g. Harrison & Feldman (2009), Wilder et al. (2011), Wilder, Feldman, & Singh (2015), Wilder, Feldman, & Singh (2015)). (See Fig. 7).

It is natural to express shape similarity using shape skeletons, because the likelihood $p(x|S)$ quantifies the degree to which the shape x fits a given model (skeleton) S, which in turn can be used to quantify how well two shapes fit a common model and the entailed part decomposition. A natural measure of shape similarity that captures this idea is the *cross-likelihood*, introduced in Briscoe (2008) (see Feldman et al. (2013), for a synopsis). The cross-likelihood is the average log-likelihood of two shapes given each other's MAP skeletons,

$$\text{crosslike}(x, y) = [-\log p(x|\text{MAP}_y) - \log p(y|\text{MAP}_x)]/2 \tag{2}$$

where MAP_x and MAP_y are the MAP skeletons for shape x and y respectively (Fig. 6). (As is conventional with likelihood measures, we take negative logs to give them the convenient form of Description Lengths which generally have the form $-\log p$.) The cross-likelihood expresses the fit of each shape to the *other* shape's best model, and thus captures the degree to which they could plausibly have arisen from the same model. (This idea is loosely inspired by the probabilistic approach to similarity introduced by Kemp, Bernstein, & Tenenbaum (2005) and Tenenbaum & Griffiths (2001).) As discussed in Briscoe (2008), the cross-likelihood is a convenient approximation to a more complete and computationally complex similarity measure, the development of which is beyond the scope of the current paper. The cross-likelihood is a reasonable approximation as long as the shapes' part structures are not radically different, and is sufficient for the current experiments.

To fit the similarity model to our data, we measured the cross-likelihood (Eq. (2)) at each point α in the morph space between the shape at α and a shape at $\alpha + \Delta\alpha$ (set at $\Delta\alpha = .1$ in the analyses). This comparison is an estimate of the local rate of representational change at each point in shape space—loosely speaking, the local derivative of the mental shape manifold, which will be larger where the shape representation is changing more rapidly, corresponding hypothetically to greater sensitivity to shape differences, and lower where the shape representation is changing more slowly, corresponding to lower sensitivity. The computed similarity at point α was then regressed against the sensitivity measured at the same point (i.e. the values plotted in Figs. 3–5). The only degrees of freedom in these fits are the two DFs associated with the linear regression between cross-likelihood and measured sensitivity, an arbitrary mapping necessary to convert units. The cross likelihood measure itself has no free parameters, other than those of the underlying skeleton estimation procedure, which we set to default values.[1] We compared the goodness-of-fit of the cross-likelihood model to the (null) constant model (indicating no variation in sensitivity with shape), and corrected for the difference in the number of degrees of freedom in the two models to yield an AIC.

Table 2 gives AICs (corrected log likelihood) of the fits of the cross-likelihood similarity measure to our discrimination data, relative to the null model. The rightmost column in the table indicates the implied likelihood ratio by which the data favors the cross-likelihood relative to the null model (with AIC differences converted to corrected log likelihoods as described in Burnham & Anderson (2002)). According to Jeffreys' well-known translation of likelihood ratios to categories of "significance," 8 of the 17 conditions provide "decisive" evidence in favor of our model, 3 "strong" evidence, 1 "substantial" evidence, and 5 provided evidence (1 in favor, 4 against) that is weak enough to be "not worth mentioning." In other words, most conditions strongly or very strongly supported our model, a few provided weaker evidence, and none clearly supported the null model.

We conclude that the deviations from uniformity in our sensitivity

[1] Matlab code for skeleton estimation can be found at http://ruccs.rutgers.edu/images/ShapeToolbox1.0.zip

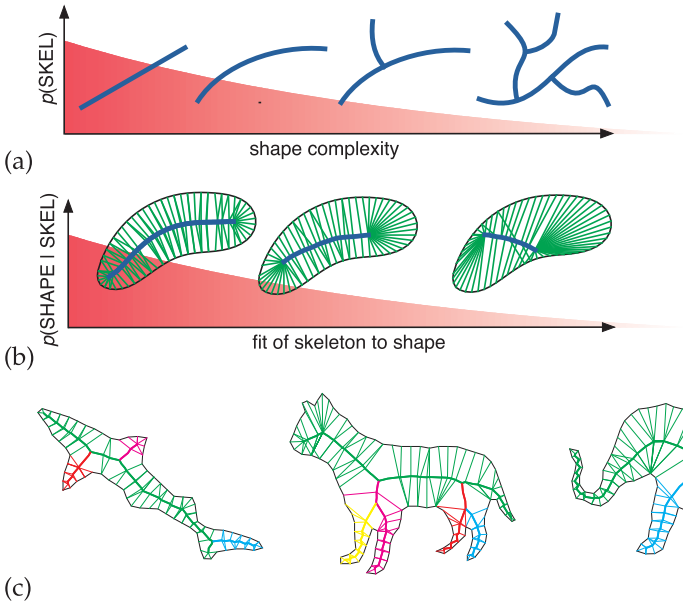

Fig. 6. Illustration of (a) prior and (b) likelihood in the Feldman and Singh (2006) shape skeleton estimation framework. The MAP skeleton (c) maximizes the product of the prior and likelihood, yielding the skeleton most likely to have generated the observed shape. Distinct colors indicate distinct axes in the axis hierarchy. Axes are shown with "ribs" (correspondences by which axial points explain contour points).

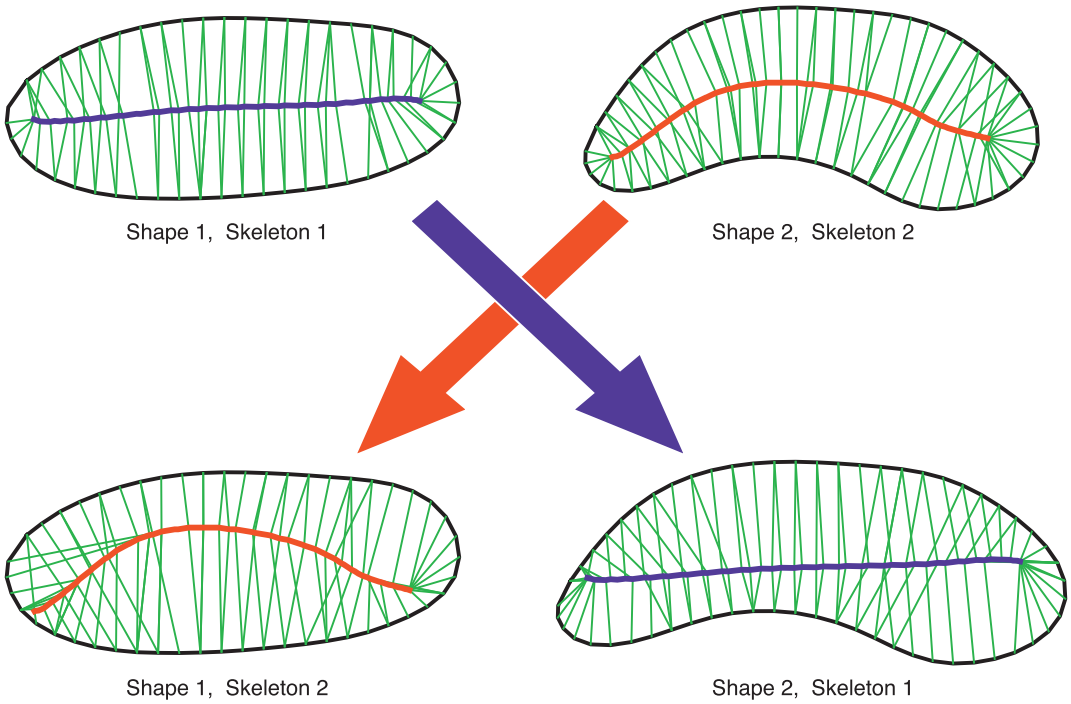

Fig. 7. Illustration of the cross-likelihood, which quantifies the average fit (likelihood) of each shape's MAP skeleton (top row, showing skeletons and ribs) to the other shape (bottom row). Likelihoods are computed as in Feldman and Singh (2006).

plots (a) are statistically robust, and (b) can be explained reasonably well by our skeletal shape similarity measure. The more different two shapes' skeletons are, and the more poorly they fit each others' best shape model, the more detectable is the difference between them.

The predictive power of our shape similarity metric stems, in our view, from its origins in the shape skeleton estimation framework. As discussed above, the likelihoods comprised in the cross-likelihood express the degree to which each shape fits particular shape models. Gradations in likelihood entail variations in fit to the model, and shape changes that alter the relative likelihood of two distinct shape models signify relatively categorical changes between distinct generative interpretations of the shape. That is, the qualitative nature of skeletal models imposes a soft probabilistic categorical structure on the underlying shape space. This implicit categorical structure includes shape factors well-known to be important to human shape representation, such as symmetry, but also encompasses categorical shape changes that have proven more difficult to quantify, like part structure, and aggregates all these factors together in a unified probabilistic measure. Hence the success of the cross-likelihood in explaining our discrimination data provides empirical support not only for the cross-likelihood similarity measure itself, but also for the entire shape skeleton estimation framework from which it is derived.

4. Conclusions

In this study, we asked participants to discriminate pairs of shapes, and found substantial variations in discriminability that cannot be explained in terms of the degree of simple geometrical difference between

Table 2
Model fits, showing AICs for the cross-likelihood similarity models, the constant model, and the implied correctd likelihood ratio favoring the former model (LR = $\exp([AIC_1 - AIC_2]/2)$; see Burnham and Anderson (2002)). Larger ratios indicate stronger evidence for the similarity model, and ratios less than 1 indicate evidence in favor of the constant model.

Condition	Skeletal model AIC	Constant model AIC	Corrected LR (skeletal/constant)
One-Part/Two-Part Set A			
No normalization	184.3472	197.1168	592.8
Normalized perimeter	392.2722	390.3364	.38
Random scale	233.1251	231.7254	.50
One-Part/Two-Part Set B			
Random scale	321.81	333.8235	406.2
Two-Part/Three-Part Set A			
Normalized area	291.0397	290.5649	.79
Normalized perimeter	602.0987	603.2586	1.8
Two-Part/Three-Part Set B			
Normalized area	133.4905	169.9436	8.24×10^{7}
Normalized perimeter	270.2524	323.2172	3.17×10^{11}
Two-Part/Three-Part Set C			
Random scale	130.4771	135.8006	14.3
Two-Part/Three-Part Set D			
Random scale	182.5402	224.903	1.58×10^{9}
Two-Part/Three-Part regular			
No normalization	231.6787	230.3486	.52
Circle/Ellipse			
Normalized area	322.0451	352.2819	3.68×10^{8}
Normalized perimeter	56.44244	96.79211	5.78×10^{8}
Random scale	143.446	225.3872	6.21×10^{17}
Ellipse/Peanut			
Normalized area	176.7095	180.053	5.32
Ellipse/Bent Ellipse			
Normalized area	240.4451	245.1462	10.49
Normalized perimeter	190.0998	196.2669	21.84

the shapes. Some contour changes induce relatively large subjective differences, and thus better discrimination, while other changes of nominally equal magnitude yield relatively small differences and poorer discrimination. Our stimulus sets included a number of theoretically interesting morph spectra, and the resulting data are very revealing about the structure of mental shape spaces. Broadly speaking, the pattern of discriminability differences were well fit by a simple skeleton-based similarity measure, the skeleton cross-likelihood, which arises from the Bayesian account of shape representation Feldman and Singh (2006). Shape pairs that were more similar according to this measure were generally less discriminable, while more dissimilar pairs were more discriminable.

We readily acknowledge that it is difficult to generalize about the precise pattern of discrimination variation in our data. In some cases discrimination is superior near highly symmetric forms such as the circle (Figs. 5a-c), while in other cases discrimination seems to be superior at an intermediate point between two subjective shape classes Fig. 5(e,f). However, regardless of the individual characteristics of particular morphs, the overall pattern is that whatever discriminability differences arise seem to be well-predicted by the skeletal similarity measure. The model is not perfect, of course, and the data show substantial noise. However the problem of shape similarity is notoriously subtle and the model fits are generally good, *especially* because no simple or easily-stated general shape principle seems to account for the variations in discrimination. Instead, discrimination appears to follow the predictions provided by the cross-likelihood similarity measure, even when those predictions are difficult to articulate through intuitive explanations. The model appears to capture some subtle and non-obvious aspects of shape discrimination.

In spite of this success, we also acknowledge that the cross-likelihood is at best a coarse similarity measure, and can almost certainly be improved upon. In a recent paper, (Erdogan & Jacobs, 2017) provide an impressively comprehensive comparison of shape classification techniques, in which the cross-likelihood measure does not predict human classification responses (in a very different set of stimuli) as well as another Bayesian model. However their model does not provide a shape similarity measure, only a classification, and cannot be readily applied to our discrimination data. Existing skeletal similarity measures based on graph matching (e.g. Goh, 2008; Sebastian & Kimia, 2005) also cannot readily explain our results, because they are based solely on skeletal topology, unlike ours which takes into account probabilistic fits of shapes to skeletons. Such methods even in principle cannot account for variations in discrimination when the skeleton is topologically constant, as in several of our experiments.

It should also be noted that our experiments so far have been limited to the 2D case, which allowed us to focus purely on shape and avoid issues such an viewpoint and self-occlusion. However, we acknowledge that a complete theory of human shape representation and similarity must include the 3D case. Although such an extension is outside the scope of the work described here, we believe that an exploration of 3D shape is an important next step as we build towards a more complete model of shape representation.

As mentioned above, the cross-likelihood can be seen as a simple approximation to a more principled probabilistic similarity model, in that it quantifies the probability each shape having arisen from the other's generative process. Loosely speaking, this approach is reasonable as long as the two shapes are not too different, but will probably break down with radically different skeletons. This makes the cross-likelihood appropriate for the discrimination tasks, such as the one reported here, in which shape pairs are necessarily very similar to each other, but it cannot be expected to generalize well to more divergent shapes. In future work, we hope to develop and test a more comprehensive similarity measure suitable for comparing shapes with arbitrarily different skeletons.

Acknowledgments

Preparation of this article was supported by NIH (NEI) Grant EY021494 to J.F. and M.S. We are grateful to Melchi Michel for helpful discussions.

References

Amir, O., Biederman, I., & Hayworth, K. J. (2012). Sensitivity to nonaccidental properties across various shape dimensions. *Vision Research, 62*, 35–43.
Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review, 61*, 183–193.
Barenholtz, E., Cohen, E., Feldman, J., & Singh, M. (2003). Detection of change in shape: An advantage for concavities. *Cognition, 89*(1), 1–9.
Biederman, I. (1987). Recognition by components: A theory of human image understanding. *Psychological Review, 94*, 115–147.
Binford, T. (1981). Inferring surfaces from images. *Artificial Intelligence, 17*, 205–244.
Blum, H. (1973). Biological shape and visual science (part I). *Journal of Theoretical Biology, 38*, 205–287.
Blum, H., & Nagel, R. N. (1978). Shape description using weighted symmetric axis features. *Pattern Recognition, 10*, 167–180.
Briscoe, E. (2008). *Shape skeletons and shape similarity (Unpublished doctoral dissertation)*. Rutgers University.
Burnham, K. P., & Anderson, D. R. (2002). *Model selection and multi-model inference: A practical information-theoretic approach.* New York: Springer.
Chen, M.-T., & Chen, K. C. (1987). A group model of form recognition under plane similarity transformations. *Journal of Mathematical Psychology, 31*, 321–337.
Cohen, E., Barenholtz, E., Singh, M., & Feldman, J. (2005). What change detection tells us about the visual representation of shape. *Journal of Vision, 5*(4), 313–321.
Cortese, J. M., & Dyre, B. P. (1996). Perceptual similarity of shapes generated from fourier descriptors. *Journal of Experimental Psychology: Human Perception and Performance,*

22(1), 133–143.
de Winter, J., & Wagemans, J. (2006). Segmentation of object outlines into parts: A large-scale integrative study. *Cognition, 99*(3), 275–325.
Denisova, K., Feldman, J., Su, X., & Singh, M. (2016). Investigating shape representation using sensitivity to part- and axis-based transformations. *Vision Research, 126*, 347–361.
Edelman, S. (1998). Representation is representation of similarities. *Behavioral and Brain Science, 21*(4), 449–467.
Erdogan, G., & Jacobs, R. A. (2017). Visual shape perception as Bayesian inference of 3D object-centered shape representations. *Psychological Review, 124*(6), 740–761.
Feldman, J. (1997). The structure of perceptual categories. *Journal of Mathematical Psychology, 41*, 145–170.
Feldman, J. (2007). Formation of visual objects in the early computation of spatial relations. *Perception & Psychophysics, 69*(5), 816–827.
Feldman, J., & Richards, W. A. (1998). Mapping the mental space of rectangles. *Perception, 27*, 1191–1202.
Feldman, J., & Singh, M. (2006). Bayesian estimation of the shape skeleton. *Proceedings of the National Academy of Sciences, 103*(47), 18014–18019.
Feldman, J., Singh, M., Briscoe, E., Froyen, V., Kim, S., & Wilder, J. D. (2013). An integrated Bayesian approach to shape representation and perceptual organization. In S. Dickinson, & Z. Pizlo (Eds.). *Shape perception in human and computer vision: An interdisciplinary perspective*. Springer.
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), 1–7.
Folstein, J. R., Palmeri, T. J., Van Gulick, A. E., & Gauthier, I. (2015). Category learning stretches neural representations in visual cortex. *Current Directions in Psychological Science, 24*(1), 17–23.
Gillebert, C. R., Op de Beeck, H. P., Panis, S., & Wagemans, J. (2009). Subordinate categorization enhances the neural selectivity in human object-selective cortex for fine shape differences. *Journal of Cognitive Neuroscience, 21*(6), 1054–1064.
Goh, W.-B. (2008). Strategies for shape matching using skeletons. *Computer Vision and Image Understanding, 110*(3), 326–345.
Goldstone, R. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General, 123*(2), 178–200.
Goldstone, R., & Hendrickson, A. T. (2010). Categorical perception. *Wiley Interdisciplinary Reviews: Cognitive Science, 1*(1), 69–78.
Harnad, S. (1987). *Categorical perception: The groundwork of cognition.* Cambridge: Cambridge University Press.
Harrison, S., & Feldman, J. (2009). Influence of shape and medial axis structure on texture perception. *JOV, 9*(6), 1–21.
Hartendorp, M. O., Van der Stigchel, S., Wagemans, J., Klugkist, I., & Postma, A. (2012). The activation of alternative response candidates: when do doubts kick in? *Acta Psychologica, 139*(1), 38–45.
Hoffman, D. D., & Richards, W. A. (1984). Parts of recognition. *Cognition, 18*, 65–96.
Hoffman, D. D., & Singh, M. (1997). Salience of visual parts. *Cognition, 63*, 29–78.
Katz, R. A., & Pizer, S. M. (2003). Untangling the Blum medial axis transform. *International Journal of Computer Vision, 55*(2/3), 139–153.
Kayaert, G., Biederman, I., & Vogels, R. (2003). Shape tuning in macaque inferior temporal cortex. *Journal of Neuroscience, 23*(7), 3016–3027.
Kemp, C., Bernstein, A., & Tenenbaum, J. B. (2005). A generative theory of similarity. *Proceedings of the 27th annual conference of the cognitive science society*.
Kendall, D. G. (1989). A survey of the statistical theory of shape. *Statistical Science, 4*(2), 87–120.
Kimia, B. B. (2003). One the role of medial geometry in human vision. *Journal of Physiology Paris, 97*, 155–190.
Kingdom, F. A. A., & Prins, N. (2010). *Psychophysics: A practical introduction.* London: Academic Press.
Kovács, I., Fehér, A., & Julesz, B. (1970). Medial-point description of shape: A representation for action coding and its psychophysical correlates. *Vision Research, 38*, 2323–2333.
Kubilius, J., Wagemans, J., & de Beeck, H. P. O. (2014). Encoding of configural regularity in the human visual system. *Journal of Vision, 14*(9) 11–11.
Landau, B., Smith, L. B., & Jones, S. S. (1988). The importance of shape in early lexical learning. *Cognitive Development, 3*, 299–321.
Liu, Z., & Kersten, D. (2003). Three-dimensional symmetric shapes are discriminated more efficiently than asymmetric ones. *Journal of the Optical Society of America A, 20*(7), 1331–1340.
Lowe, D. G. (1987). Three-dimensional object recognition from single two-dimensional images. *Artificial Intelligence, 31*, 355–395.
Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information.* San Francisco: Freeman.
Newell, F. N., & Bulthoff, H. H. (2002). Categorical perception of familiar objects. *Cognition, 85*(2), 113–143.
Notman, L. A., Sowden, P. T., & Ozgen, E. (2005). The nature of learned categorical perception effects: a psychophysical approach. *Cognition, 95*(2), 1–14.
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. *Journal of Neuroscience, 28*(40), 10111–10123.
Palmer, S. E., & Hemenway, K. (1978). Orientation and symmetry: Effects of multiple, rotational, and near symmetries. *Journal of Experimental Psychology: Human Perception and Performance, 4*(4), 691–702.
Panis, S., Vangeneugden, J., & Wagemans, J. (2008). Similarity, typicality, and category-level matching of morphed outlines of everyday objects. *Perception, 37*(12), 1822–1849.
Richards, W., Dawson, B., & Whittington, D. (1988). Encoding contour shape by curvature extrema. *Natural computation.* Cambridge, MA: M.I.T. Press.
Rosch, E. H. (1973). Natural categories. *Cognitive Psychology, 4*, 328–350.
Sebastian, T. B., & Kimia, B. B. (2005). Curves vs. skeletons in object recognition. *Signal Processing, 85*, 247–263.
Shepard, R. N. (1957). Stimulus and response generalization: a stochastic model relating generalization to distance in psychological space. *Psychometrika, 22*(4), 325–345.
Siddiqi, K., Shokoufandeh, A., Dickinson, S., & Zucker, S. (1999). Shock graphs and shape matching. *International Journal of Computer Vision, 30*, 1–24.
Siddiqi, K., Shokoufandeh, A., Dickinson, S. J., & Zucker, S. W. (1998). Shock graphs and shape matching. *Proceedings of the sixth international conference on computer vision* (pp. 222). IEEE Computer Society.
Siddiqi, K., Tresness, K. J., & Kimia, B. B. (1996). Parts of visual form: Psychophysical aspects. *Perception, 25*, 399–424.
Sigala, N., & Logothetis, N. K. (2002). Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature, 415*(6869), 318–320.
Singh, M., & Hoffman, D. D. (2001). Part-based representations of visual shape and implications for visual cognition. In T. Shipley, & P. Kellman (Vol. Eds.), *From fragments to objects: Segmentation and grouping in vision, advances in psychology: vol. 130*, (pp. 401–459). New York: Elsevier.
Singh, M., Seyranian, G. D., & Hoffman, D. D. (1999). Parsing silhouettes: The short-cut rule. *Perception & Psychophysics, 61*(4), 636–660.
Sprote, P., & Fleming, R. W. (2016). Bent out of shape: The visual inference of non-rigid shape transformations applied to objects. *Vision Research, 126*, 330–346.
Tenenbaum, J. B., & Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences, 24*(4), 629–640.
Wagemans, J. (1992). Perceptual use of non-accidental properties. *Canadian Journal of Psychology, 46*(2), 236–279.
Wagemans, J. (1995). Detection of visual symmetries. *Spatial Vision, 9*(1), 9–32.
Wagemans, J., Van Gool, L., Swinnen, V., & Van Horebeek, J. (1993). Higher-order structure in regularity detection. *Vision Research, 33*(8), 1067–1088.
Wilder, J., Feldman, J., & Singh, M. (2011). Superordinate shape classification using natural shape statistics. *Cognition, 119*, 325–340.
Wilder, J., Feldman, J., & Singh, M. (2015). Contour complexity and contour detection. *Journal of Vision, 15*(6), 1–16.
Wilder, J., Feldman, J., & Singh, M. (2015). The role of shape complexity in the detection of closed contours. *Vision Research*.
Witkin, A. P. (1983). Scale-space filtering. *Proceedings of the eighth international joint conference on artificial intelligence: vol. 2*, (pp. 1019–1022). San Francisco, CA, USA: Morgan Kaufmann Publishers Inc Retrieved fromhttp://dl.acm.org/citation.cfm?id=1623516.1623607.
Yamane, Y., Carlson, E. T., Bowman, K. C., Wang, Z., & Connor, C. E. (2008). A neural code for three-dimensional object shape in macaque inferotemporal cortex. *Nature Neuroscience, 11*(11), 1352–1360.