Perceived orientation of complex shape reflects graded part decomposition

Elias H. Cohen
Department of Psychology and Center for Cognitive Science, Rutgers University, New Brunswick, NJ, USA

Manish Singh
Department of Psychology and Center for Cognitive Science, Rutgers University, New Brunswick, NJ, USA

Although the orientation of line segments and simple shapes is a well-studied area of vision, little is known about geometric factors that influence perceived orientation of complex multipart shapes. The study of these factors is of interest because it allows for an insight into the basic problem of how local geometric attributes are integrated perceptually into a global shape representation. We examined the perceived orientation of two-part shapes using an adjustment method and a 2AFC task. In particular, we investigated the influence of the perceptual salience, or distinctiveness, of a part—as defined by the turning angles at its boundaries—and its area relative to the main "base" part. In contrast to previous results on simple shapes, our results exhibited large and systematic deviations of perceived orientation from the principal axis of the shape. For shapes with sharp part boundaries, perceived global orientation deviated maximally from the principal axis and was approximated by the axis of the main base part of the shape. With weakening part boundaries, the perceived orientation gradually approached the principal axis of the entire shape, reflecting that both parts were taken into account in estimating orientation. The results are consistent with a differentially weighted principal-axis computation in which the attached part is given systematically lower weighting with increasing turning angles at the part boundaries. They thus allow a quantitative characterization of part salience in terms of *part independence*: Turning angles at a part's boundaries determine the extent to which its influence is perceptually separable from the rest of the shape. We suggest that Robust Statistics may provide a useful framework for quantifying the influence of part segmentation on visual estimation.

Keywords: shape representation, parts, components, structural descriptions, orientation, principal axis

Introduction

Initial measurements by biological visual systems are confined to small, local "apertures" in retinal images. Observers' perception of object structure, on the other hand, is known to be strongly influenced by nonlocal attributes of an object's geometry. A basic problem in vision science is to understand how local image measurements are integrated into the global representation of object shape. A natural way to approach shape integration is to investigate how local properties of a shape's geometry influence the perception of a global shape attribute. In this study, we use orientation as the global attribute to be estimated. In particular, we investigate how local geometric determinants of a part's visual salience affect the perceived overall orientation of two-part shapes.

Part-based representation of shape

Prior research indicates that the human visual system represents complex shapes in terms of simpler parts, rather than as unstructured templates. It decomposes shapes into smaller semi-independent parts and organizes shape representation in terms of these parts and their spatial relationships (see, e.g., Biederman, 1987; Hoffman & Richards, 1984; Marr & Nishihara, 1978; Palmer, 1977; Singh, Seyranian, & Hoffman, 1999). Such structured representations[1] have the benefit of allowing for greater robustness under changes in viewing conditions—for instance, involving changes in articulated pose—because the representation of an object's shape can be dissociated from any particular spatial configuration that its parts may take.

Hoffman and Richards (1984) proposed that the visual system uses negative minima of curvature—points of locally maximal magnitude of curvature in concave regions along a shape's occluding contour—to define candidate boundaries between parts. This *minima rule* was motivated by the *principle of transversality*, according to which the intersection of two smooth surfaces almost surely (i.e., with probability 1) generates a concave tangent discontinuity at their locus of intersection[2] (Guillemin & Pollack, 1974).

Part segmentation at negative minima has indeed been found to explain a number of phenomena in shape perception (see Singh & Hoffman, 2001, for a review), including the perception of symmetry and repetition (Baylis & Driver, 1994), changes in perceived shape associated with reversals of figure and ground (Driver & Baylis, 1996; Hoffman & Singh, 1997), the perception of

doi: 10.1167/6.8.4 Received July 31, 2005; published July 14, 2006 ISSN 1534-7362

transparency (Singh & Hoffman, 1998), the localization of vertex height (Bertamini, 2001), visual search asymmetries (Hulleman, te Winkel, & Boselie, 2000; Wolfe & Bennett, 1997; Xu & Singh, 2002), and differential performance in comparing two probes along a shape's outline (Barenholtz & Feldman, 2003). Additionally, change detection involving complex shapes indicates a heightened sensitivity to concavities (Barenholtz, Cohen, Feldman, & Singh, 2003; Cohen, Barenholtz, Singh, & Feldman, 2005). More recently, we have shown using a segment-identification task (Cohen & Singh, 2004) that observers are substantially better at determining whether or not a given contour segment has been taken from the outline of a given shape if it is segmented at negative minima of curvature, rather than at positive maxima or at inflections.

Recent work in physiology also provides evidence for curvature-based coding of contours in area V4 of the macaque. Different sets of V4 neurons have been found to be selective for contour segments with different magnitudes of curvature and, moreover, for either positive or negative curvature (i.e., locally convex or concave segments; see Pasupathy & Connor, 1999, 2001). Selectivity for sign of curvature is also found in the inferotemporal cortex, where cell responses have been shown to generalize over mirror reversals and contrast reversals of shapes but not over figure-ground reversals—which necessarily switch the sign of curvature at each point along a contour (i.e., reverse the roles of local convexities and concavities; Baylis & Driver, 2001).

Part salience

The studies cited above have brought us closer to an understanding of the geometric properties that the visual system uses in segmenting shapes into parts. However, our understanding of how the representation of global shape is organized by human vision using segmented parts is still rather limited. In computational vision, hierarchical representational schemes involving tree structures have been proposed, where branches of the shape's medial-axis "skeleton" (Blum, 1973) corresponding to different parts are organized in a tree hierarchy, with smaller parts nested within larger ones (e.g., Kimia, Tannenbaum, & Zucker, 1995; Rom & Medioni, 1993). Although these representational schemes have a great deal of intuitive appeal for biological vision as well, the precise way in which the human visual system organizes the global representation of shape using parts has yet to be determined.

It is clear that not all parts have the same status in the visual representation of a shape: A toe, for instance, has quite a different status in the structural representation of the human form than the torso. However, even with sheer size and level in a shape's hierarchy equated, two parts can have very different perceptual salience. The smaller attached part

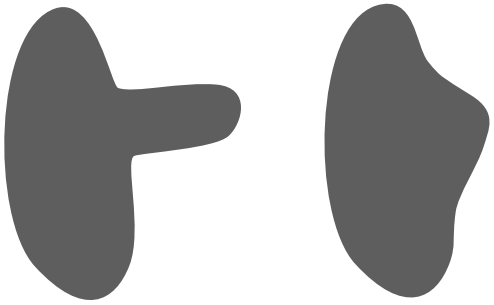

Figure 1. Demonstration of the difference in the perceptual salience of a part arising from the turning angles at its part boundaries and its protrusion (ratio of part perimeter to cut length).

in each of the two shapes in Figure 1 has the same size and is at the same level in its shape's hierarchy (here, a simple one, consisting of a single small part protruding out of a larger "base"). However, the part on the left is considerably more *salient* perceptually than the one on the right. This example illustrates two geometric factors proposed by Hoffman and Singh (1997) in their theory of part salience: *part-boundary strength* and *part protrusion*. The part on the left is delineated by negative minima that have substantially greater turning angles than the one on the right,[3] and its protrusion—defined by the ratio of its perimeter to the length of its part cut—is also substantially higher.

Previous work has documented some influences of the turning angle, curvature, or both at part boundaries on part salience—for example, via its influence on figure-ground assignment (Hoffman & Singh, 1997), performance on a task comparing two probes along a shape's bounding contour (Barenholtz & Feldman, 2003), and likelihood of choosing a negative minimum as a part boundary (De Winter & Wagemans, 2006).

In a recent series of experiments, we directly tested the influence of turning angle at negative minima of curvature (as well as other geometric properties) on the perceptual salience of a part, using a 4AFC task (see Figure 2). We employed the experimental rationale—first used by Palmer (1977) to measure part goodness—that observers should be systematically better at verifying/identifying parts with higher perceptual salience. On each trial, a randomly generated 2D shape was briefly flashed and masked. Following this, a contour segment taken from the bounding contour of the shape was presented, followed by a screen divided into four quadrants (Figure 2a). The observers' task was to indicate the location of the contour segment on the original shape, by choosing one of the four quadrants. The results showed that observers' accuracy on this task increases steadily with the turning angle at part boundaries—thereby indicating that parts defined by higher turning angles at their part boundaries are

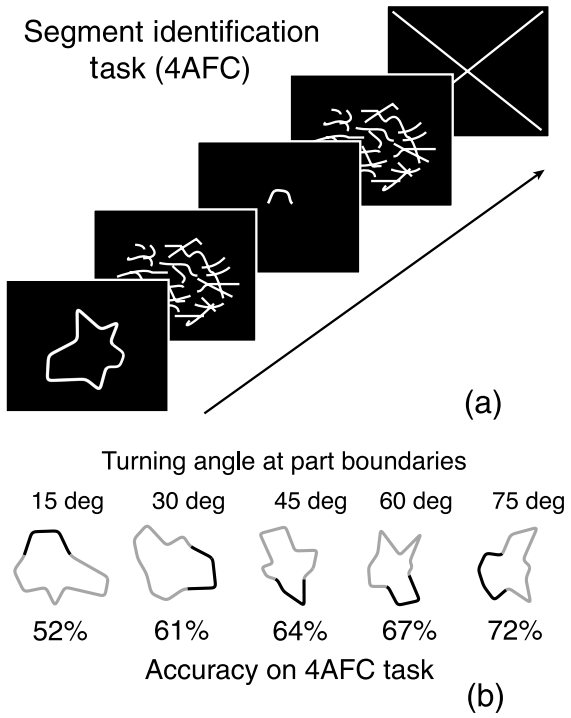

Figure 2. Illustration of the method and results from Cohen and Singh (2004). (a) The trial sequence. Observers identified the location of a contour segment on a previously presented shape by choosing one of the four screen quadrants. (b) Results showed that accuracy in the 4AFC task increases systematically with the turning angles at the part boundaries. (The results also demonstrated a systematic influence of segment length and part protrusion—not shown here).

perceptually more salient and more readily identified as perceptual "units" (Figure 2b).

Part salience as independence

The fact that parts within the same level of a shape's hierarchy can have very different perceptual salience implies that a discrete representational format such as a hierarchical tree must be supplemented with graded information. A common way of thinking about part salience is in terms of the *representational strength* of a part in a shape's description: The higher the part's salience is, the stronger is its representation (e.g., Hoffman & Singh, 1997; Rom & Medioni, 1993; Siddiqi & Kimia, 1995). Such a view is consistent with claims in the literature that perceptually salient parts enjoy special status, in terms of being remembered, attended, or named more readily (e.g., Bower & Glass, 1976; Hoffman & Singh, 1997; Palmer, 1977; Reed, 1974).

A complementary way of viewing part salience is in terms of *part independence*: Highly salient parts have a higher likelihood of being segmented and are, thus, more readily perceived as independent units. Under the part-independence interpretation, part salience may be construed as the extent to which a part's representation is perceptually separable from the rest of shape. Parts with sharper turning angles at their boundaries, for instance, are more likely to be perceptually segmented and represented as distinct units.

Part independence is a complementary aspect of part salience—one that highlights a somewhat different consequence. Whereas previous work has generally emphasized that more salient parts receive increased processing resources (in terms of attention, memory, etc.), part salience may have a somewhat different effect in contexts involving the visual estimation of a *global* property of a complex shape (i.e., when integration over the entire shape is required). A highly salient part is likely to be perceived as more separate from the rest of the shape. Hence, it is likely to be weighted *less strongly* in the estimation of a global property because its contribution is less compatible with the bulk of the shape.

These considerations suggest the hypothesis that a part's salience should systematically influence observers' ability to integrate the part into a global shape estimate: The higher a part's salience is, the more *weakly* the part's contribution should be incorporated into the global estimate. In the experiments reported here, we tested this hypothesis in the context of visual estimation of shape orientation.

Perception of orientation

Orientation is one of the basic summary descriptors of an object's geometry. Apart from location and size, the orientation of an object in the environment is one of its most salient attributes and one that is crucial for interaction and manipulation (such as determining grip position or anticipating motion). Whereas the localization of an object is determined based on its first-order statistics—the mean of the coordinates of its interior points (namely, its center of mass[4])—the orientation is determined by its second-order statistics—that is, the "spread" of its points around the center of mass. Specifically, it is computed based on the direction of the first principal axis (i.e., the eigenvector of the covariance matrix corresponding to the largest eigenvalue). This is equivalent to the direction of the line that minimizes the sum of the squared (perpendicular) distances to all points that constitute the interior of the shape.[5]

The principal axis has been shown to provide a perceptually valid measure of orientation for simple shapes and dot clusters, in the context of both perception (e.g., Boutsen & Marendaz, 2001; Lansky, Yakimoff, Radil, & Mitrani, 1989; Li & Westheimer, 1997) and object manipulation (e.g., Turvey, Burton, Pagano, Solomon, & Runeson, 1992). Studies using dot clusters have shown that the principal axis (which corresponds to their least

squares line; see footnote 5) consistently predicts their perceived orientation (Lansky, Yakimoff, & Radil, 1987; Lansky et al. 1989). The precision with which observers estimate cluster orientation was found to depend on the correlation strength around the least squares line: The stronger the correlation, the more precisely observers identified that direction. Nevertheless, observers exhibited a high degree of accuracy in estimating the orientation (i.e., without systematic bias from the principal axis) even for weakly defined axes (e.g., with correlation coefficients <.23).

Li and Westheimer (1997) showed that observers are sensitive to the *implicit orientation*[6] of simple symmetric shapes (such as ellipses or the letter "X") and that estimates of implicit orientation exhibit many of the same characteristics as the explicit orientation of a line segment. For instance, implicit orientation was shown to produce an oblique effect (see also Liu, Dijkstra, & Oomes, 2002). Orientation was perceived with the greatest speed and accuracy when the shape was vertical or horizontal, despite the presence of oblique edges within the shape. (Note that the implicit orientation of an ellipse, or the letter "X," corresponds precisely to the direction of its principal axis.) Based on their results, Li and Westheimer suggested that the automatic computation of global shape orientation may be performed by mechanisms closely related to those computing the explicit orientation of a line segment. Consistent with this suggestion, Boutsen and Marendaz (2001) have shown, using a visual-search paradigm, that the principal axis of a complex shape (with a clear dominant axis but no salient part structure) is computed sufficiently early to produce pop-out effects based on orientation differences.

Given that cells in early visual cortical areas exhibit orientation tuning for local edges, a particularly simple neural model of shape orientation could be based simply on pooling all local orientation measurements along a shape's bounding contour. Burbeck and Zauberman (1997) examined this hypothesis by studying orientation discrimination of rectangular shapes whose longer sides were modulated sinusoidally. Bias in determining the shapes' orientation was found to depend on the relative phase of the sinusoidal modulations on the two sides and, moreover, for the in-phase modulations, on their spatial frequency as well. Burbeck and Zauberman reasoned that edge-based models cannot account for this influence of relative phase because, in their stimuli, the mean edge orientation was unaffected by their phase manipulation. Moreover, they argued that the influence of edge-modulation frequency also rules out a model based on the responses of large-scale units whose receptive fields encompass the entire shape: If a receptive field is large enough to cover the entire shape, changes in the frequency of modulation of its bounding contour should not affect its response. Based on these results, Burbeck and Zauberman proposed that the medial axis—instantiated in terms of the *core model* (see Burbeck & Pizer, 1995)—must play a role in the perception of object orientation. (Their manipulation of relative phase and spatial frequency had a systematic influence on the medial axis). Although Burbeck and Zauberman did not consider such a model, the computation of principal axis also predicts an influence of relative phase and frequency modulation in their shapes. (Appendix A shows the results of the principal-axis computation applied to their shapes.)

Oomes and Dijkstra (2002) investigated the perceived orientation of ellipsoids and other simple three-dimensional (3D) shapes, presented stereoscopically. Observers adjusted the orientation of a 3D cross—comprising three mutually perpendicular lines—to match the perceived orientation of a 3D shape. Observers' settings were found to be precise, but in many cases, they exhibited large deviations (up to 25 deg) from the 3D principal axis. By decomposing observers' 3D orientation settings into slant (orientation in depth) and tilt (orientation in the image plane), Oomes and Dijkstra found that the errors arose almost entirely from the settings of slant. Observers' orientation estimates were thus strongly biased toward the principal axis of the projected *silhouette* of the 3D objects.

The above studies indicate that human vision is proficient at using the principal axis (at least in the image plane) to compute overall orientation for a variety of stimuli. Previous studies have, however, mostly used relatively simple shapes that lack salient part structure. In particular, the influence of clearly demarcated component parts on the perceived orientation of complex multipart shapes has not been directly investigated. In this study, we begin this investigation by focusing on the context of two-part shapes in which the parts are arranged in a simple hierarchical relationship: a small part protruding out from a larger base (see Figure 3). We address the questions of how the global orientation of such shapes is visually determined and with what degree of precision it is represented.

There are two natural hypotheses concerning how global orientation might be computed for shapes with multiple parts. The first—*homogenous computation hypothesis*—postulates that all points within a shape are treated uniformly in computing the principal axis, irrespective of part structure. In other words, the computation of principal axis proceeds just as it does for simpler shapes. According to the second—*part-based computation hypothesis*—the computation of shape orientation explicitly takes into account the decomposition of the shape into parts. Principal axes may be computed separately for individual parts and then integrated into a global orientation estimate or the principal-axis computation may proceed by assigning different weights to different points within the shape, depending on which part they belong to.

In the following experiments, we tested the "null" hypothesis of the homogenous computation of orientation in the context of two-part shapes: a small part protruding out of a dominant base. We investigated how the decomposition of a shape into parts—in particular, the perceptual

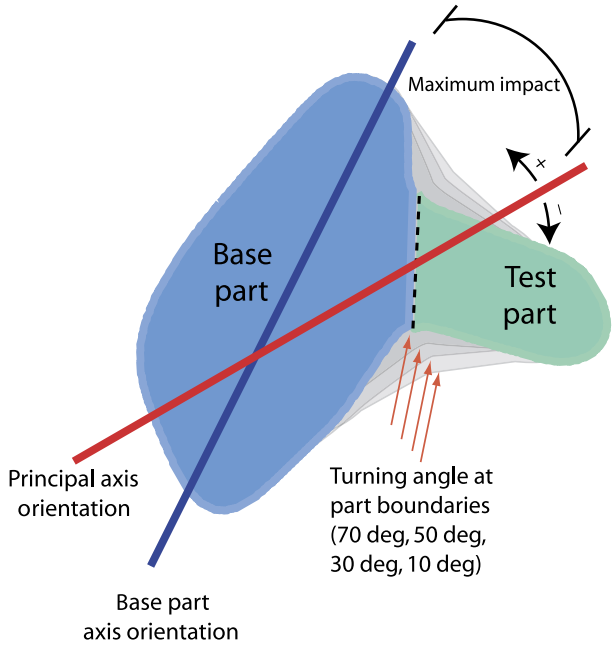

Figure 3. Demonstration of the construction of the test shapes. Shapes were generated using two ellipses with random perturbations added to their boundaries—a base part and an attached part. Part-boundary strength was varied by manipulating the turning angles at the part boundaries. The "maximal impact" was defined as the difference between the shape's principal axis orientation and the orientation of the base part. Observer settings were measured relative to the principal axis. Angular deviations in the direction of the base part were designated as positive.

salience of the attached part—influences the perceived overall orientation of the shape. To this effect, we manipulated two geometric properties of the two-part shapes: the size of the smaller (attached) part and the turning angles at its part boundaries. Perceived shape orientation was measured using a method of adjustment (Experiment 1) and a 2AFC task (Experiment 2).

Experiment 1

Method

Observers

Sixteen undergraduate students at Rutgers University served as naive observers, in exchange for course credit.

Stimuli

Test shapes were constructed using two overlapping ellipses of different sizes, a base part and an attached part (see Figure 3). The size of the base part was fixed (major axis = 5.34 deg of visual angle; minor axis = 2.67 deg of visual angle). The attached part had the same aspect ratio as the base part; its major axis could take one of four lengths: 19%, 23%, 27%, or 31% of the major axis of the larger part. The ellipse corresponding to the attached part was centered on the bounding contour of the larger ellipse, at a polar angle of 45, 135, 225, or 315 deg, measured from the center of the larger ellipse. The smaller ellipse was oriented orthogonally to the larger one. Multiple variants of the composite shape were created by introducing random perturbations to points on the boundary of each ellipse.

The bounding contour of each ellipse was sampled at 12 points, uniformly spaced in polar angle. The location of each point was then perturbed using an angular perturbation drawn from a uniform distribution on [−5 deg, 5 deg], and a radial perturbation drawn from a uniform distribution on $[-r/15, r/15]$, where r is the radial distance of the point from the center of the ellipse.

Part-boundary strength was manipulated by varying the magnitude of turning angles at the negative minima of curvature. These could take four values: 10, 30, 50, and 70 deg. The turning angles were varied by shifting each negative minimum along the bisector line of the angle formed by its closest neighbors. Finally, a smoothing operation was applied to the bounding contours by fine sampling and convolving with a 1D Gaussian filter (see, e.g., Mokhtarian and Mackworth, 1992). The smoothed shape was filled in with uniform gray (18 cd/m^2) and presented against a black background (0.03 cd/m^2). Sixteen instances were generated for each of the 16 shape types, yielding a total of 256 test shapes. Figure 4 shows an instance of a shape for each of the 4 × 4 combinations of turning angle and part size.

The test shapes were interleaved with randomly generated shapes, which comprised one third of the stimuli. The random shapes were generated as 11-point polygons and then smoothed by fine sampling and convolving with a 1D Gaussian. The location of each vertex was selected by first projecting 11 equally spaced radial axes from the center of the screen, with radial distance drawn from a uniform distribution on [1.07, 7.12] deg of visual angle. A random angular offset was then applied to each vertex (uniformly distributed on [−5.5 deg, 5.5 deg]). The random shapes contained between one and five parts. The data from the random shapes were not analyzed. Their purpose was to prevent overlearning of the test stimuli and discourage the development of specific strategies for the test shapes.

Design and procedure

Two independent variables were manipulated: *part size* and *turning angle* at the part boundaries, each with four levels (see Stimuli section). Thus, 16 shape types were possible.

The experiment was divided into blocks of 24 trials each. Each observer ran 16 experimental blocks, for a total of 384 adjustments. These were preceded by 2 practice blocks.

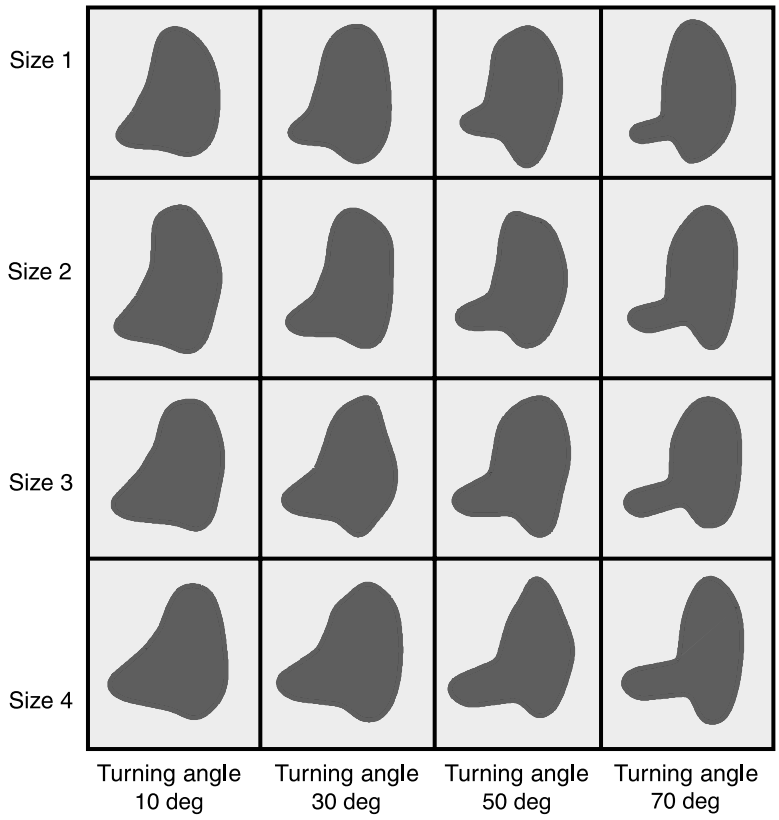

Figure 4. Stimulus design in Experiment 1: For each of the 4 × 4 combinations of turning angle and part size, 16 different shapes were generated, with small random variations (see text for details). This figure shows one shape instance for each combination of the two variables.

On each trial, the observer was presented with the following sequence: (1) a study shape for 1,000 ms, (2) a mask for 500 ms, and (3) an adjustable probe line. The observers' task was to adjust the orientation of the probe line using a track ball. They were instructed as follows: "Rotate the red probe line to match the perceived orientation of the entire shape. If you see multiple possible orientations, choose the strongest one. There is no right or wrong answer." Shapes were presented at random orientations, rotated about their center of mass. They were presented centered on the screen (again, with respect to their center of mass). After the presentation of the shape and mask, the red probe line was displayed at the center of the screen. Its initial orientation was randomly determined on each trial.

Results

Each orientation setting was encoded in terms of its angular deviation from the principal axis of the shape. Because the attached part could occupy one of four positions with respect to the base part and because the entire shape could be presented at any orientation, the raw orientation data were first transformed into a canonical coordinate framework. Deviations from the principal axis toward the base-part orientation were designated positive; those toward the smaller, attached part were designated negative (see Figure 3). Initial data analysis revealed no systematic influence of shape handedness (the position of the attached part relative to the base part) on orientation settings. The data were thus collapsed over left-handed and right-handed shapes.

Figure 5a plots the mean orientation settings, expressed as angular deviations from the principal axis, with error bars depicting standard errors.[7] The orientation settings exhibited a significant dependence on turning angle at the part boundaries, $F(3, 45) = 42.29$, $p < .0001$. As the turning angles at the part boundaries increase, observers' orientation settings deviate increasingly from the shape's principal axis, in the direction of the base-part axis. The settings also exhibited a systematic dependence on part size, $F(3, 45) = 21.35$, $p < .0001$. The larger the attached part is, the greater is the deviation of perceived orientation from the shape's principal axis.

Additionally, the interaction between part-boundary strength and part size was also significant, $F(9, 135) = 8.73$, $p < .0001$. As is evident in Figure 5a, the influence of

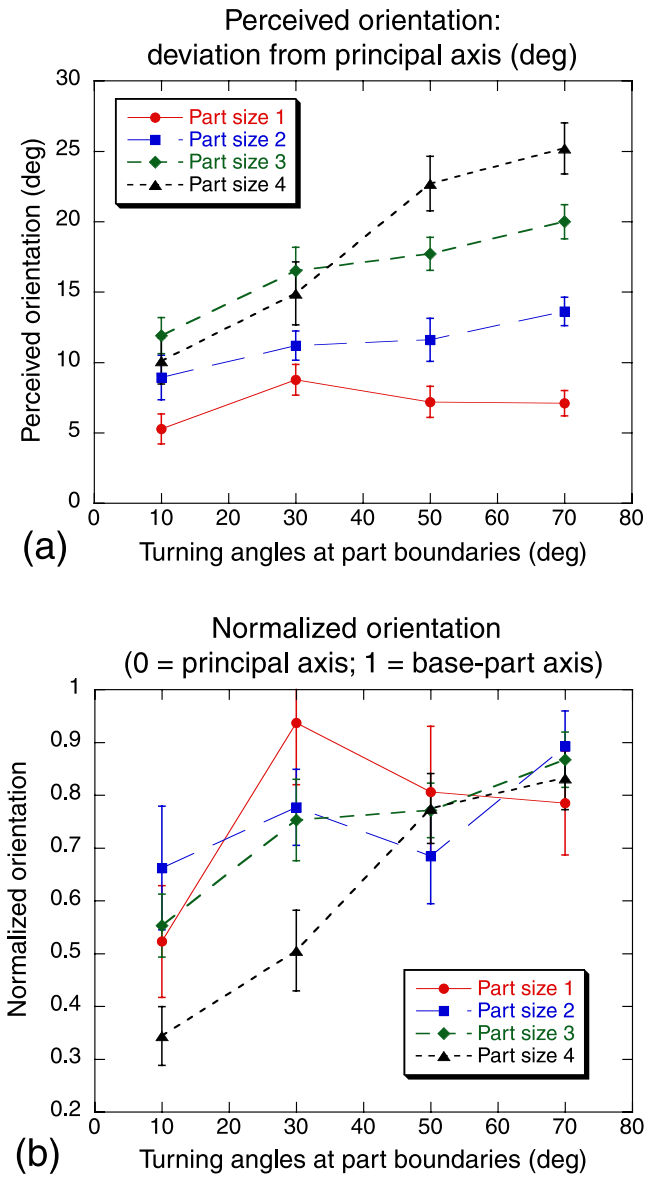

Figure 5. (a) Mean orientation settings in Experiment 1, measured as angular deviations from the principal axis, plotted as a function of boundary turning angle and part size. Shifts in the positive direction correspond to shifts in perceived orientation toward the base-part axis. Error bars indicate standard errors of the mean. (b) Mean orientation settings normalized by the maximal impact of the part—that is, its geometric influence on the orientation of the principal axis. In the normalized plots, 0 on the *y*-axis corresponds to the principal-axis orientation and 1 corresponds to the base-part orientation.

turning angle on observers' settings is considerably greater for larger parts. (See Appendix B for polar histograms of observers' orientation settings, for each of the 16 conditions.)

Although both turning angle and part size had a significant effect on observers' orientation settings, whether these effects constitute genuine perceptual influences of these variables on orientation estimates is not evident from the above tests. It is possible, in particular, that these effects simply reflect the *geometric* influence that these variables have on the principal-axis orientation—in particular, on the angular separation between the principal axis and the base-part axis. As part size increases, for instance, the principal-axis orientation deviates increasingly from the base-part orientation. Thus, if observers' orientation settings corresponded to the orientation of the base part, for instance, their settings expressed as deviations from the principal axis would exhibit a systematic increase with part size. This would clearly not constitute a genuine perceptual influence of part size on shape orientation, however, beyond simply its geometric influence on the shape's principal axis.

To test whether turning angle and part size have a genuine *perceptual* influence, we normalized the orientation settings in Figure 5a by the respective angular separations between principal-axis and base-part orientation (i.e., by the *geometric* influence that the presence of the part has on the shape's principal axis, or "maximal impact"). Figure 5b plots the data from Experiment 1 in terms of this normalized orientation.

Along the "normalized orientation" scale on the *y*-axis, 0 corresponds to the orientation of the principal axis and 1 to the orientation of the base-part axis.[8] When the analyses were performed on normalized orientation, the influence of part size became marginal, $F(3, 45) = 2.84$, $p = .0485$. Thus, part size does not exert a reliable influence on perceived orientation, beyond its geometric influence on the principal axis. The influence of turning angle, however, continued to be highly significant, $F(3, 45) = 10.469$, $p < .0001$. This indicates that turning angle at the part boundaries has a genuine perceptual influence on orientation estimates, beyond any geometric influence.[9] Finally, the interaction between turning angle and part size was significant, $F(9, 135) = 2.83$, $p < .005$. The influence of turning angle on normalized orientation tends to be greater for larger parts.

To examine trends in setting variability, we performed tests of heteroscedasticity on individual observers' data. These tests revealed a significant influence of part size on setting variance for 11 of the 16 observers and a significant influence of turning angle for 6 observers. Increase in part size tended to increase setting variance, whereas increase in boundary turning angle tended to decrease setting variance.

Discussion

The results of Experiment 1 demonstrate that turning angles at part boundaries exert a strong influence on the perceived orientation of two-part shapes. With increase in turning angles, orientation settings deviate increasingly from the principal axis and approach the axis of the base

part. These results are consistent with the interpretation of part salience as independence: Increase in the strength of the part boundaries increases the percept of a more separate/independent perceptual unit, and the part is then weighted more weakly in the overall computation of orientation—or, loosely speaking, "ignored" to a greater extent. (We will return to this point in greater detail in the General discussion section.) This is also consistent with the decreasing trend in setting variability associated with increasing boundary strength: The computation of orientation is more precise when it is driven mostly by the base part because this is simply an ellipse with small random perturbations. With increase in part size, orientation settings were pulled more toward the attached part. However, part size cannot be said to have a genuine perceptual influence: Its effect on perceived orientation largely disappeared once its geometric influence on the principal axis was factored out. The increase in setting variability with increasing part size is to be expected for a number of reasons, including an increase in total shape area and a decrease in overall elongation (defined by the ratio of the eigenvalues associated with the two eigenvectors). The ratio of the two eigenvalues (lower to higher) increased systematically, becoming closer to 1, with increasing part size—the mean values for the four part sizes being .36, .40, .43, and .45. (A ratio of 1 corresponds to a shape having no dominant axis.)[10]

Experiment 2

This experiment used a 2AFC task to further investigate the influence of turning angles at part boundaries on perceived orientation. Observers viewed two-part shapes presented at random orientations followed by a probe line. Their task was to indicate the direction of the probe's angular offset (clockwise or counterclockwise) relative to the perceived orientation of the shape. The point of subjective equality (PSE; 50% threshold) was used to estimate the perceived orientation, and the slope of a Weibull fit was used to estimate precision. Part size was fixed at the largest size used in Experiment 1. A turning angle of 0 deg was added to the design, yielding five levels of turning angle. Five hundred distinct instances were generated for each value of turning angle, for a total of 2,500 shapes.

Methods

Observers

Four experienced observers with normal or corrected-to-normal visual acuity participated in the experiment. Three were naive to the purpose of the experiment; the fourth was author E.C. (observer O2).

Stimuli

The generation scheme for the test shapes and the parameters of the base part were identical to Experiment 1. The size of the attached part was fixed at the largest value used in Experiment 1 (the length of the major axis of the smaller ellipse was 31% of the major axis of the base ellipse). Five values of turning angle were used. Four were the same as the values used in Experiment 1 (10, 30, 50, and 70 deg). An additional value of 0 deg of turning angle was included (which smoothed out the boundaries between the two parts). Five hundred instances of test shapes were generated for each level of turning angle. Each shape was presented only once to an observer. Only the two-part test shapes were used (random shapes were not included).

Design and procedure

On each trial, the observer was presented with the following sequence: (1) a test shape for 500 ms, (2) a mask for 250 ms, and (3) a probe line presented until response (see Figure 6). Observers' 2AFC task was to indicate the direction of the probe's angular offset relative to the shape—that is, in which direction (clockwise or counterclockwise) the probe should be rotated to align with the perceived orientation of the shape.

Probe lines were presented at 1 of 10 orientations relative to the principal axis of the test shape: ranging from −30 to 51 deg, in steps of 9 deg. (As before, positive angles refer to rotations toward the base part of the shape, whereas negative angles refer to rotations toward the smaller attached part; see Figure 3). These angular offsets were chosen to encompass the entire range of orientations between the principal axis (by convention, 0 deg) and the axis of the base part (mean orientation = 30.28 deg). Each combination of turning angle (5 values) and angular offset from principal axis (10 values) was repeated in 50 trials, yielding a total of 2,500 experimental trials per observer.

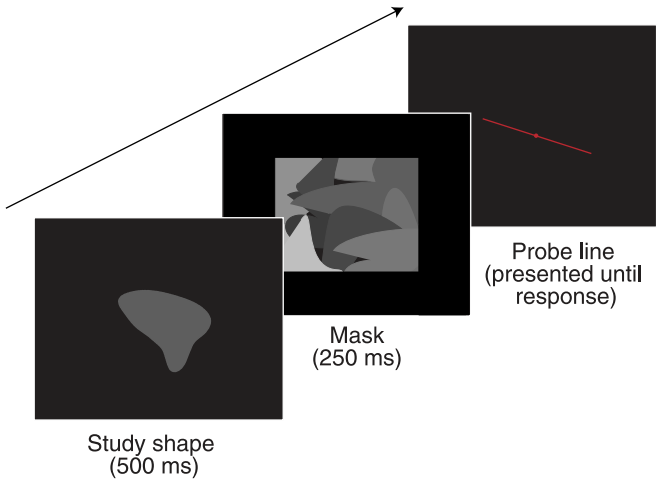

Figure 6. The trial sequence for Experiment 2.

Each observer participated in five experimental sessions, preceded by one practice session. Each session contained 500 trials.

Results and discussion

Figure 7 shows the psychometric fits to individual observers' data for each turning-angle condition. Weibull functions were fit using the *psignifit* software, version 2.5.6 for Matlab (see Wichmann & Hill, 2001). The PSE (50% threshold) was used to measure perceived orientation, and the slope of the Weibull fit was used to measure the precision of the orientation estimate.

All observers exhibited a significant increase in PSE with turning angles at the boundaries (Figure 8a). As the turning angles increased, perceived orientation shifted systematically away from the principal axis (by convention, 0 deg) and approached that of the base part (mean = 30.28 deg). As in Experiment 1, these estimates of perceived orientation were normalized by their respective geometric impacts (i.e., the angular separation between the base-part axis and principal axis; see Figure 8b). Turning angle continued to have a significant effect, even with the geometric influence of the part factored out. This again indicates that the observed effect of turning angle constitutes a genuine perceptual influence on shape orientation—one that is not reducible to its geometric influence. (Note that the normalized PSEs are shown only for the nonzero values of turning angle because, when turning angles are zero, the difference between base-part and principal-axis orientation is equal to zero, and thus, normalized PSEs are undefined.) As before, this shift toward the base-part axis suggests that, as the smaller part becomes more distinctive, it is treated as more independent and separable and weighted more weakly in the computation of orientation.[11]

In addition to the shift in PSE, turning angle also influenced the precision of the orientation estimates. With increase in turning angle at the part boundaries, the slope of the Weibull fit exhibited an increasing trend for three of the observers—indicating a greater precision for shapes with more distinctive parts (see Figure 9). The increase in precision is naturally seen as reflecting the fact that, as the attached part becomes more distinctive, the computation of orientation is driven to a greater extent by the base part, which has a well-defined axis of elongation.

General discussion

The results of both experiments revealed a systematic influence of part salience on the perceived orientation of two-part shapes. For sharply defined part boundaries (large turning angles), perceived orientation deviated maximally from the principal axis and was approximated instead by the axis of the base part—suggesting that the

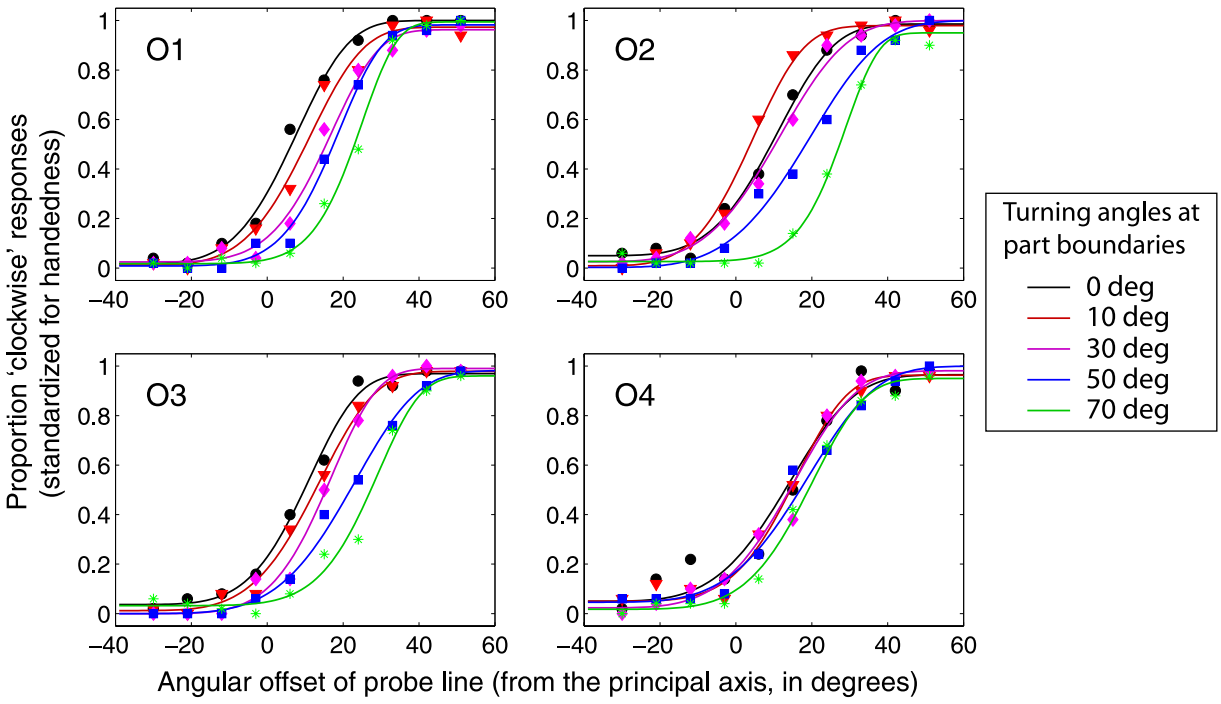

Figure 7. Weibull fits to the 2AFC data for each boundary-strength condition in Experiment 2, shown for the four observers. The thresholds and slopes derived from these fits are shown in Figures 8 and 9.

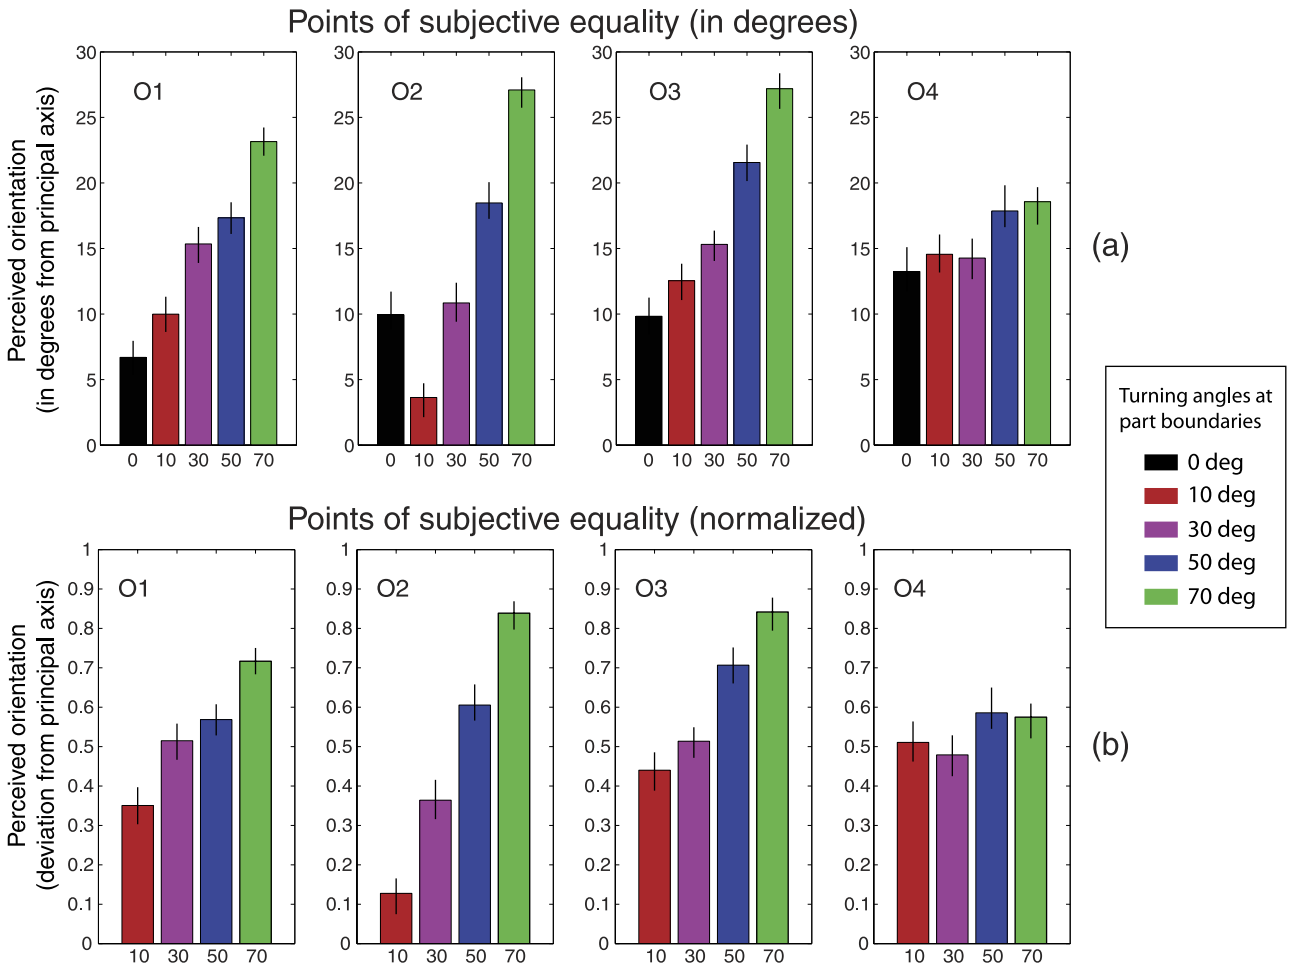

Figure 8. (a) PSEs obtained from the 50% thresholds of the psychometric fits. These indicate the perceived orientation for each turning-angle condition, measured as angular deviation (in degrees) from the overall principal axis. (b) The PSEs normalized by the geometric impact of the part for each respective turning-angle condition. On the normalized scale, 0 indicates principal-axis orientation and 1 indicates base-part orientation. (The normalized PSEs are shown only for the nonzero values of turning angle because they are undefined for shapes without part structure.) Error bars indicate 95% confidence intervals.

attached part is largely ignored in computing orientation in these cases. With decreasing strength of the part boundaries (shallower turning angles), the perceived orientation gradually approached the direction of the principal axis, reflecting that both parts were being taken into account. Normalization of the orientation data showed that this influence of turning angle constitutes a genuine perceptual effect; the effect of turning angle remained highly significant even when its geometric influence on the principal axis was factored out. Consistent with this pattern of bias, the precision also tended to be highest for sharply defined boundaries—when only the base part was utilized.

These results thus argue against the homogenous computation hypothesis for the visual estimation of orientation. The visual system does not treat all points within a shape uniformly in computing its orientation. Rather, it explicitly takes into account the structural decomposition of the shape into parts and treats points within the two parts differently. We have suggested that part salience may be naturally interpreted in terms of perceptual independence. A highly salient part is perceptually more independent (or separable) from the rest of the shape; hence, its contribution is weighted more weakly in the visual estimation of the shape orientation. A low-salience part is less separable (i.e., more integrated into the shape), and its influence is therefore more on an equal footing with the rest of the shape.

Differentially weighted principal-axis computation

A natural way to quantify the observed influence of part salience is in terms of a differential weighting assigned to

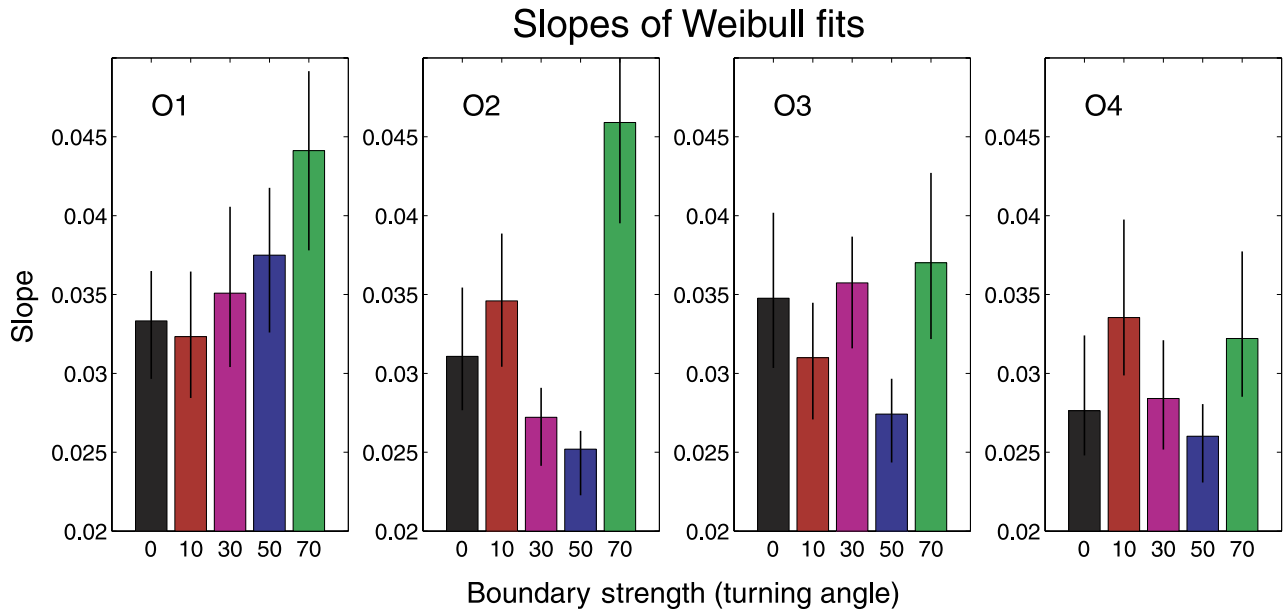

Figure 9. Slopes of the Weibull fits for each boundary-strength condition and observer. Error bars indicate 95% confidence intervals.

the attached part, in the principal-axis computation. In other words, the principal axis may be computed with points within the base part assigned a "full" weight of 1, and points within the attached part assigned a partial weight of w (where $0 \leq w \leq 1$). When $w = 1$, points within the attached part are on an equal footing with the base part and, as a result, the computed orientation coincides simply with that of the standard principal axis of the entire shape (see Figure 10a).

When $w = 0$, the attached part is entirely deweighted and not used at all (or ignored) in the principal-axis computation; the computed orientation thus coincides with the principal axis of the base part alone. As the weight w decreases from 1 to 0, the computed orientation shifts gradually from the principal axis of the entire shape to that of the base part alone (see Figure 10a).

With such a computation in mind, one may recharacterize the orientation data in terms of the relative weight of the attached part that would generate the observed orientation for any given shape. Figure 10b shows the result of such a recoding of the Experiment 1 data. For each of the 256 experimental shapes used in Experiment 1, we computed the value of relative weight w for the attached part, which, when used in a differentially weighted principal-axis computation for that shape, would generate the observed orientation. We then collapsed these computed weights across all shape instances within a given condition (i.e., a combination of turning angle and part size).

The resulting plot (Figure 10b) looks similar to the normalized orientation plot in Figure 5b, although it should be noted that the scales of the y-axis in the two cases have different meanings. The normalized orientation data in Figure 5b indicate where the observed orientation falls on the unit scale between the principal-axis orientation and base-part orientation, whereas the values on the y-axis of Figure 10b indicate the differential weight to points within the attached part, necessary in a principal-axis computation, to generate that observed orientation.[12] Nevertheless, the basic conclusion remains unchanged: Turning angle at the part boundaries has a genuine perceptual influence on perceived orientation that is not reducible to a geometric influence on the principal axis.

Part decomposition and Robust Statistics

The above scheme involving the differential weighting of distinct parts within a shape, to estimate its overall orientation, is analogous to the *Robust Statistics* approach to statistical estimation. This branch of statistics aims to develop estimators that are robust against violations of one's assumptions concerning the underlying model (e.g., the data points being sampled independently from a single Gaussian distribution). The methods developed allow for principled ways to deal with the presence of outliers—data points that either involve gross errors (e.g., in entry or coding) or were otherwise generated from a different process than the one under study.

Treating such points uniformly with the rest of the data can lead to large errors in estimation (linear regression, for instance, is notorious for its sensitivity to outliers). Although outliers are sometimes identified and removed by visual inspection, this approach is

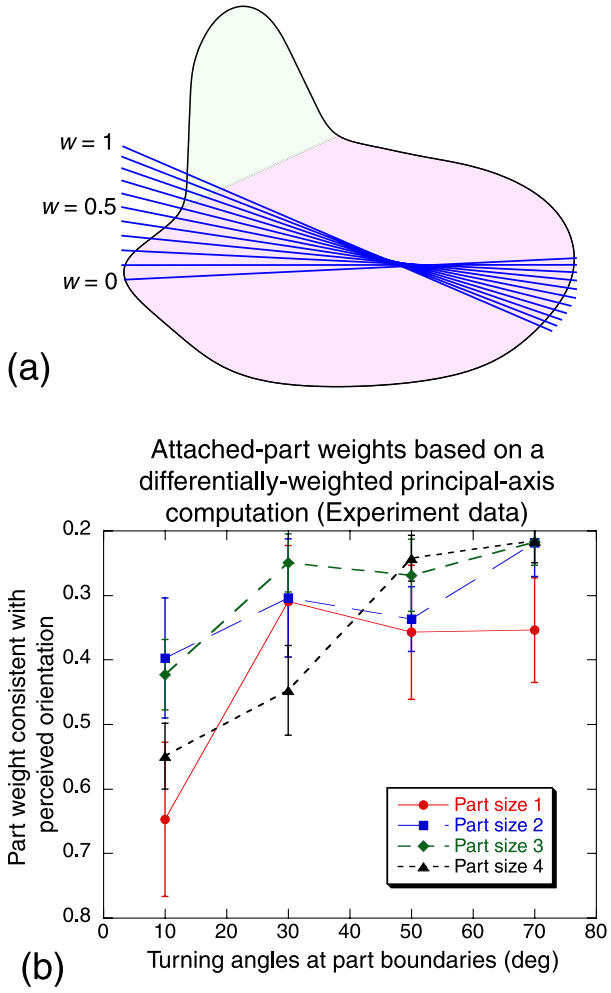

Figure 10. (a) Demonstration of the result of a differentially weighted principal-axis computation. Points within the small attached part are assigned partial weights, *w,* ranging from 0 to 1, relative to those within the base part. When $w = 1$, the attached part is on an equal footing with the base part, and the computed orientation coincides with the standard principal axis. As w approaches 0, the computed orientation shifts systematically toward the orientation of the base part. (b) The data from Experiment 1 recoded in terms of the relative weight, *w,* of the attached part necessary to generate the observed orientation for any given shape.

problematic because it can force researchers to make hard decisions early on, based on ill-defined criteria. Robust methods adopt the more principled approach of weighting data points differently based on some quantity that reflects (implicitly or explicitly) the likelihood that the data arose from the same generative process as the bulk of the data. Thus, rather than making a binary choice between including or excluding a potential outlier, its contribution is down-weighted (ignored) to varying extents, depending on how likely it is that the data point arose from a different generative process (see, e.g., Hampel, 1974; Landy, Maloney, Johnston, & Young, 1995).

What is the significance of Robust Statistics for visual processing? Any estimation of a visual property (e.g., location, orientation, motion, or depth) requires pooling information over an extended image region. However, such pooling is meaningful only if it is performed over an appropriately defined region that includes precisely those points that correspond to the relevant object or surface being estimated. Mixing in the motion signals from two different surfaces, for instance, will generally lead to a poor motion estimate for both (e.g., McDermott & Adelson, 2004). Successful visual estimation thus depends critically on perceptual segmentation of the image into appropriate units—within which information may be meaningfully combined.

Although, from a physical point of view, objects and surfaces may be treated as essentially discrete entities, the same is not true of perception. Given the ambiguity and uncertainty that the visual system faces in computing object and surface boundaries from image data (arising from the multiplicity of interpretations, as well as from the noisy measurements), perceptual segmentation is more naturally regarded as probabilistic "inference" (see, e.g., Elder & Goldberg, 2002; Feldman, 2001; Geisler, Perry, Super, & Gallogly, 2001; Hon, Maloney, & Landy 1997; Yuille, Fang, Schrater, & Kersten, 2004, for the probabilistic approach applied to curve integration). Probabilities are assigned to segmentation/grouping hypotheses reflecting the degree of "confidence" that a certain image region corresponds to the same contour/surface as another. We suggest that this approach is likely to be useful in the study of perceptual part segmentation as well.

In this article, we investigated the visual segmentation of parts within a single object. Specifically, we examined how the perceptual salience, or distinctiveness, of a part within a shape influences the perceived overall orientation of the shape. We found that, as the part becomes more distinctive (via an increase in the turning angles at its boundaries), the perceived orientation of the shape deviates increasingly from its principal axis and approaches the orientation of the base part. This result is consistent with a differentially weighted computation of principal axis, in which points within the attached part are given systematically lower weights relative to the base part, as the salience of the attached part increases. This strategy suggests that a Robust Statistics approach is being employed by the visual system in the computation of object orientation: The higher the part's salience is—the greater its perceptual independence from the rest of the shape—the lower is its contribution to the orientation estimate. (The probability p reflecting the perceptual independence of a part may be quantified simply as $p = 1 - w$, where w is the relative weight assigned to it in the differentially weighted principal-

axis scheme.) Robust Statistics is thus likely to provide a useful quantitative framework within which to capture the influence of perceptual segmentation on visual estimation.

Conclusions

In investigating the influence of local geometric properties—turning angles at part boundaries—on overall perceived orientation, this study used the context of two-part shapes in which the parts are arranged in a simple hierarchical relationship: a small part protruding out of a larger base. Orientation estimates using both a method of adjustment (Experiment 1) and a 2AFC task (Experiment 2) showed that perceived orientation for these shapes is predicted by a differentially weighted principal-axis computation in which points with the smaller part are given a systematically lower weight with increasing turning angles at the part boundaries— hence, increasing perceptual distinctiveness of the part.

The shapes used in the present experiments were designed to exhibit a clear dominance hierarchy between the parts—with one part clearly seen as the dominant base. One part was noticeably larger in size, and its edges leading to the part boundaries exhibited "good continuation" with each other (consistent with the definition of "limb" by Siddiqi & Kimia, 1995). These factors led to the natural interpretation of the larger part being the base, with the smaller one branching out of it. How the visual system determines the dominance hierarchy of parts in general, however, remains an open question. In particular, how do relative size and local "branching" cues combine to determine perceived dominance hierarchy—especially when they are inconsistent with one another? How large must one part be, relative to the other, to generate a clear dominance hierarchy? For instance, if the size of the smaller part were to be gradually increased, how would the weighting assigned to the smaller part in a differentially weighted computation be affected? These and other interesting questions on geometric properties that determine the perceived dominance hierarchy of parts, as well as their influence on the visual estimation of global shape properties, await systematic investigation.

Appendix A: Principal-axis computation on shapes from Burbeck and Zauberman (1997)

This appendix briefly describes the results of principal-axis computation on Burbeck and Zauberman's (1997) shapes.

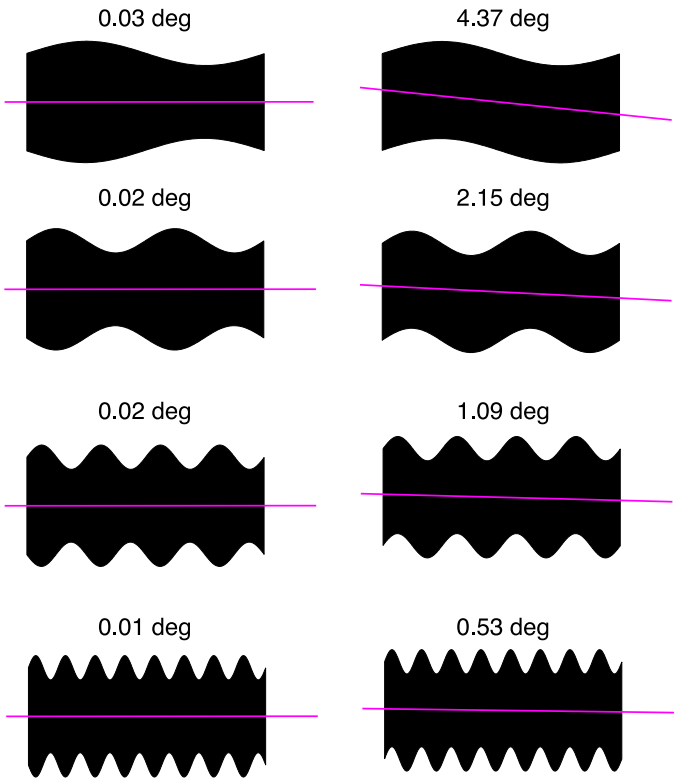

Figure A1. The orientation of the principal axis (in degrees, measured from the major axis of the underlying rectangle) for the shape manipulations used by Burbeck and Zauberman (1997). The predictions of the principal-axis computation mirror the influence of relative phase and frequency of modulation observed in their experiment.

Their shapes were generated from rectangles whose longer sides were modulated sinusoidally. The aspect ratio of the underlying rectangles was 5:2, and the amplitude of the sinusoidal modulation was 25% of the rectangles' width. The shapes varied in the frequency of the sinusoidal modulation (1, 2, 4, or 8 cycles per length of rectangle), as well as the relative phase of modulation across the two sides (in-phase or out of phase; see Figure A1). The orientation of the principal axis of each shape was measured relative to orientation of the underlying rectangle.

The results demonstrate that principal-axis computation predicts a systematic influence of relative phase on overall orientation. The deviations of the computed orientation from the underlying rectangle were consistently larger for the shapes with in-phase modulations (Figure A1; bottom row). Moreover, in the case of in-phase modulations, the computed orientation also exhibited a systematic dependence on the frequency of modulation. Both of these effects mirror the psychophysical results of Burbeck and Zauberman (1997). Thus, although the medial axis (or the core model) may well turn out to play a role in the perception of shape orientation more generally, resorting to it is not necessary to explain their specific findings.

Appendix B

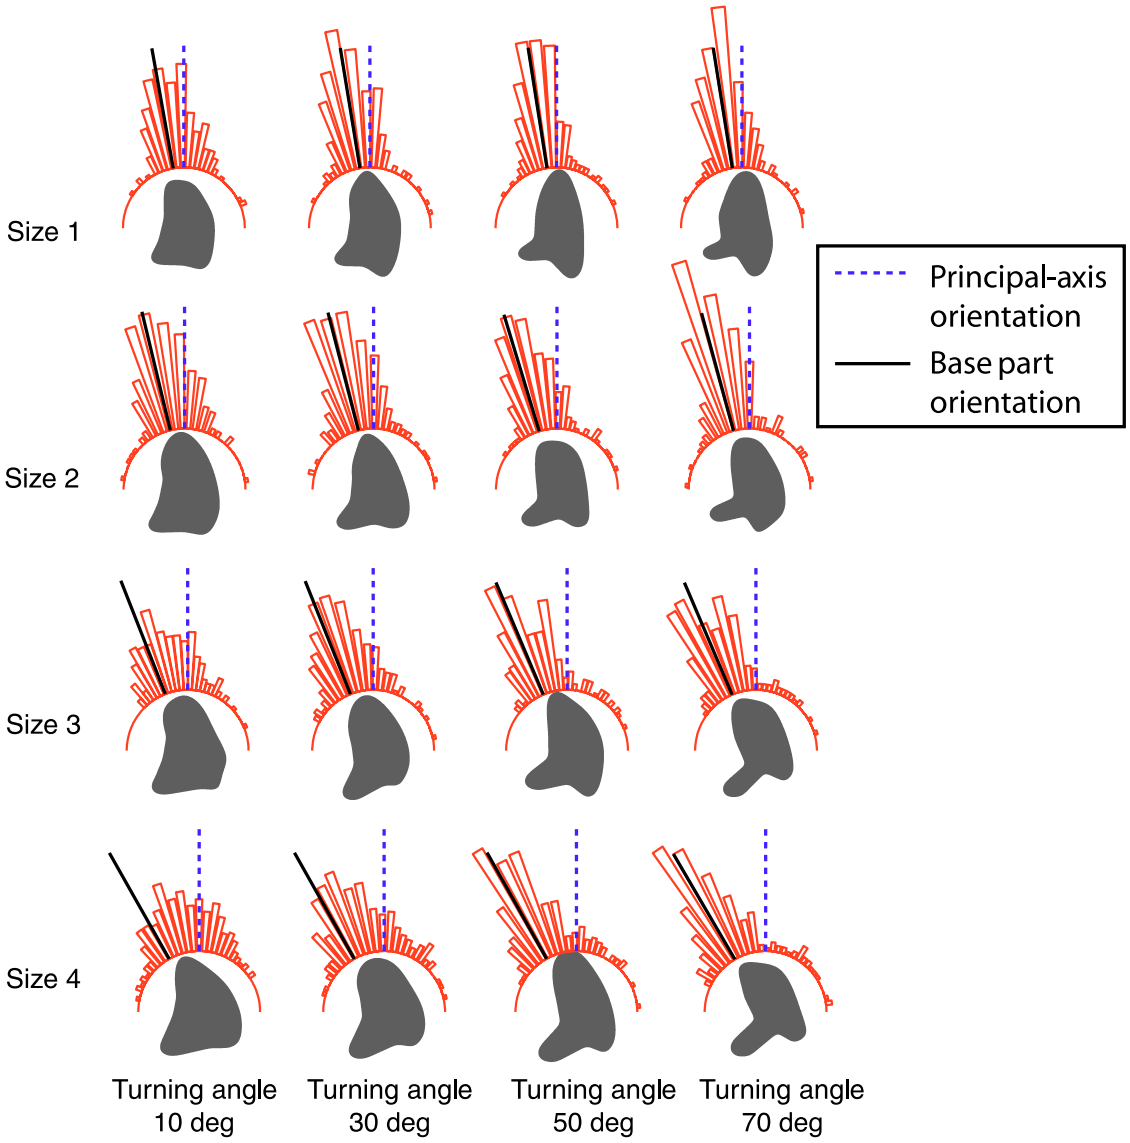

Figure B1. Pooled histograms of observer responses in Experiment 1, shown in polar coordinates. Each histogram panel includes data from all 16 shapes within a condition and shows one representative shape (of the 16 variants) from that condition. The orientations are plotted as angular deviations from the principal axis of the entire shape (the vertical dashed line). The orientation of the base part (mean for the 16 variants) is indicated by the oblique solid line. The increased tilt of the base part with increasing part size reflects the fact that larger attached parts have a greater geometric impact on the principal axis of the shape.

Acknowledgments

We are grateful to Doug DeCarlo, Jacob Feldman, Eileen Kowler, and Larry Maloney for helpful discussions, as well as to Stephen Palmer and a second, anonymous referee for comments on a previous version of the manuscript. This work was supported by NSF grant BCS-0216944.

Commercial relationships: none.
Corresponding author: Elias H. Cohen.
Email: ecohen@sunyopt.edu.
Address: 33 West 42nd Street, New York, NY 10036, USA.

Footnotes

[1]*Structural descriptions* have sometimes been taken to be synonymous with volumetric-component approaches to object recognition (e.g., Biederman, 1987; Marr & Nishihara, 1978). The notion of a structural description is considerably more general, however, and does not logically entail either that the parts be three-dimensional or that they belong to a predefined class of shape primitives (such as generalized cones).

[2]Treating part boundaries as smoothed versions of transversal intersections of surfaces leads to the 3D version of the minima rule: Part boundaries are defined by loci of negative minima of curvature along a set of lines of curvature on a smooth surface. Taking the projection onto an image plane gives the 2D version of the minima rule (see Hoffman & Richards, 1984), which we use here.

[3]For a smooth negative minimum, the part-boundary strength is determined by the normalized curvature at the point of minimum and by the turning angle within a local region containing the negative minima—defined by the nearest inflection point along the contour on each side (see Hoffman & Singh, 1997).

[4]Recent work has shown that the "averaging" performed by the visual system, to localize a shape, is not based simply on the distribution of luminance values or contrast energy but is strongly influenced by the *shape* denoted by the distribution (see Melcher & Kowler, 1999; Vishwanath & Kowler, 2003).

[5]Throughout, we will use "principal axis" to refer to the "first principal axis." It should be noted that computing the principal-axis direction is similar to *mutual regression,* in that the squared errors to be minimized are measured perpendicular to the candidate line, rather than in a fixed "vertical" direction as in standard linear regression (i.e., the squared residuals $(\hat{y} - y)^2$ of the predicted variable).

[6]By the implicit orientation of a shape, Li and Westheimer simply mean the perceived orientation that does not correspond to the orientation of any of its constituent edges. To use one of their examples, the shape of the letter "X" is perceived to be vertical although all of its local edges are oblique.

[7]Means and standard deviations for individual observers were first computed using both standard (linear) and circular statistics (see, e.g., Mardia, 1972). Values obtained with the two methods were very strongly correlated (r^2 = .996 for means, r^2 = .998 for standard deviations), which is to be expected given the small angular range within which the measurements fall. We thus report only the standard statistics throughout.

[8]It should be noted that, when expressed in normalized terms, the data for the smallest part-size condition become intrinsically more variable because of the small angular difference (mean = 9.3 deg) between the principal axis and base-part axis in this case. As a result, small variations in computed orientation due to random variations in the bounding contour of the shapes translate into large differences on the normalized scale ranging, as it does from the principal-axis orientation (0) to the base-part orientation (1).

[9]In fact, the angular separation between principal-axis orientation and base-part orientation was found to not differ significantly across the different values of turning angle.

[10]Variations in turning angle, on the other hand, did not produce a systematic change in overall elongation: The mean ratios of eigenvalues for increasing turning angles were .42, .40, .41, and .41.

[11]In postexperimental interview, Observer 4 revealed that she had been explicitly attempting to locate the principal axis for each shape (which she referred to as "the line on which the shape could be balanced"). Although effects were somewhat diminished for this observer, the influence of increasing turning angles was still evident (see Figure 8b). The fact that her strategy was specifically geared toward ignoring part structure makes the influence of part structure on her orientation estimates that much more convincing.

[12]Note that the scale in Figure 10b is reversed relative to Figure 5b because the relative weighting of the attached part is inversely related to part salience—and, hence, to the angular deviation of the orientation settings from the global principal axis. As in Figure 5b, however, the data corresponding to the smallest part size are intrinsically more variable because of the small angular difference between the principal-axis and base-part orientation in this case. In particular, the small variations in computed orientation due to random variations in the bounding contour of the shapes translate into large differences on the normalized scale.

References

Barenholtz, E., Cohen, E. H., Feldman, J., & Singh, M. (2003). Detection of change in shape: An advantage for concavities. *Cognition, 89,* 1–9. [PubMed]

Barenholtz, E., & Feldman, J. (2003). Visual comparisons within and between object parts: Evidence for a single-part superiority effect. *Vision Research, 43,* 1655–1666. [PubMed]

Baylis, G. C., & Driver, J. (1994). Parallel computation of symmetry but not repetition in single visual objects. *Visual Cognition, 1,* 377–400.

Baylis, G. C., & Driver, J. (2001). Shape-coding in IT cells generalizes over contrast and mirror reversal, but not figure-ground reversal. *Nature Neuroscience, 4,* 937–942. [PubMed]

Bertamini, M. (2001). The importance of being convex: An advantage for convexity when judging position. *Perception, 30,* 1295–1310. [PubMed]

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review, 94,* 115–147. [PubMed]

Blum, H. (1973). Biological shape and visual science: Part I. *Journal of Theoretical Biology, 38,* 205–287. [PubMed]

Boutsen, L., & Marendaz, C. (2001). Detection of shape orientation depends on salient axes of symmetry and elongation: Evidence from visual search. *Perception & Psychophysics, 63,* 404–422. [PubMed] [Article]

Bower, G. H., & Glass, A. L. (1976). Structural units and the redintegrative power of picture fragments. *Journal of Experimental Psychology: Human Learning and Memory, 2,* 456–466. [PubMed]

Burbeck, C. A., & Pizer, S. M. (1995). Object representation by cores: Identifying and representing primitive spatial regions. *Vision Research, 35,* 1917–1930. [PubMed]

Burbeck, C. A., & Zauberman, G. S. (1997). Across-object relationships in perceived object orientation. *Vision Research, 37,* 879–884. [PubMed]

Cohen, E. H., 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, http://journalofvision.org/5/4/3/, doi:10.1167/5.4.3. [PubMed] [Article]

Cohen, E. H., & Singh, M. (2004). The graded nature of parts in shape representation: Insights from a segment-identification task [Abstract]. *Journal of Vision, 4,* 663a, http://journalofvision.org/4/8/663/, doi:10.1167/4.8.663.

De Winter, J., & Wagemans, J. (2006). Segmentation of object outlines into parts: A large-scale integrative study. *Cognition, 99,* 275–325. [PubMed]

Driver, J., & Baylis, G. C. (1996). Edge-assignment and figure-ground segmentation in short-term visual matching. *Cognitive Psychology, 31,* 248–306. [PubMed]

Elder, J. H., & Goldberg, R. M. (2002). Ecological statistics of Gestalt laws for the perceptual organization of contours. *Journal of Vision, 2*(4), 324–353, http://journalofvision.org/2/4/5/, doi:10.1167.2.4.5. [PubMed] [Article]

Feldman, J. (2001). Bayesian contour integration. *Perception & Psychophysics, 63,* 1171–1182. [PubMed]

Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research, 41,* 711–724. [PubMed]

Guillemin, V., & Pollack, A. (1974). *Differential topology*. Englewood Cliffs, NJ: Prentice-Hall.

Hampel, F. R. (1974). The influence curve and its role in robust estimation. *Journal of the American Statistical Association, 69,* 383–393.

Hoffman, D. D., & Richards, W. A. (1984). Parts of recognition. *Cognition, 18,* 65–96. [PubMed]

Hoffman, D. D., & Singh, M. (1997). Salience of visual parts. *Cognition, 63,* 29–78. [PubMed]

Hon, A., Maloney, L., & Landy, M. (1997). The influence function for visual interpolation. *SPIE Conference Proceedings, 3016,* 409–419.

Hulleman, J., te Winkel, W., & Boselie, F. (2000). Concavities as basic features in visual search: Evidence from search asymmetries. *Perception & Psychophysics, 62,* 162–174. [PubMed]

Kimia, B. B., Tannenbaum, A. R., & Zucker, S. W. (1995). Shapes, shocks, and deformations: 1. The components of two-dimensional shape and the reaction-diffusion space. *International Journal of Computer Vision, 15,* 189–224.

Landy, M. S., Maloney, L. T., Johnston, E., & Young, M. (1995). Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research, 35,* 389–412. [PubMed]

Lansky, P., Yakimoff, N., & Radil, T. (1987). On visual orientation of dot patterns. *Biological Cybernetics, 56,* 389–396. [PubMed]

Lansky, P., Yakimoff, N., Radil, T., & Mitrani, L. (1989). Errors in estimating the orientation of dot patterns. *Perception, 18,* 237–242. [PubMed]

Li, W., & Westheimer, G. (1997). Human discrimination of the implicit orientation of simple symmetrical patterns. *Vision Research, 37,* 565–572. [PubMed]

Liu, B., Dijkstra, T. M., & Oomes, A. H. (2002). The beholder's share in the perception of orientation of 2D shapes. *Perception & Psychophysics, 64,* 1227–1247. [PubMed] [Article]

Mardia, K. V. (1972). *Statistics of directional data*. London: Academic Press.

Marr, D., & Nishihara, H. K. (1978). Representation and recognition of three-dimensional shapes. *Proceedings of the Royal Society of London: Series B, 200,* 269–294. [PubMed]

McDermott, J., & Adelson, E. H. (2004). The geometry of the occluding contour and its effect on motion interpretation. *Journal of Vision, 4*(10), 944–954,

http://journalofvision.org/4/10/9/, doi:10.1167/4.10.9. [PubMed] [Article]

Melcher, D., & Kowler, E. (1999). Shapes, surfaces and saccades. *Vision Research, 39,* 2929–2946. [PubMed]

Mokhtarian, F., & Mackworth, A. (1992). A theory of multiscale, curvature-based shape representation for planar curves. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 14,* 789–805.

Oomes, A. H., & Dijkstra, T. M. (2002). Object pose: Perceiving 3-d shape as sticks and slabs. *Perception & Psychophysics, 64,* 507–520. [PubMed] [Article]

Palmer, S. E. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology, 9,* 441–474.

Pasupathy, A., & Connor, C. E. (1999). Responses to contour features in macaque area V4. *Journal of Neurophysiology, 82,* 2490–2502. [PubMed] [Article]

Pasupathy, A., & Connor, C. E. (2001). Shape representation in area V4: Position-specific tuning for boundary conformation. *Journal of Neurophysiology, 86,* 2505–2519. [PubMed] [Article]

Reed, S. K. (1974). Structural descriptions and the limitations of visual images. *Memory and Cognition, 2,* 329–336.

Rom, H., & Medioni, G. (1993). Hierarchical decomposition and axial shape description. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 11,* 823–839.

Siddiqi, K., & Kimia, B. B. (1995). Parts of visual form: Computational aspects. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 17,* 239–251.

Singh, M., & Hoffman, D. D. (1998). Part boundaries alter the perception of transparency. *Psychological Science, 9,* 636–660.

Singh, M., & Hoffman, D. D. (2001). Part-based representations of visual shape and implications for visual cognition. In T. Shipley & P. Kellman (Eds.), *From fragments to objects: Segmentation and grouping invision, advances in psychology* (Vol. 130, pp. 401–459). New York: Elsevier.

Singh, M., Seyranian, G., & Hoffman, D. D. (1999). Parsing silhouettes: The short-cut rule. *Perception & Psychophysics, 61,* 636–660. [PubMed]

Turvey, M. T., Burton, G., Pagano, C. C., Solomon, H. Y., & Runeson, S. (1992). Role of the inertia tensor in perceiving object orientation by dynamic touch. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 714–727. [PubMed]

Vishwanath, D., & Kowler, E. (2003). Localization of shapes: Eye movements and perception compared. *Vision Research, 43,* 1637–1653. [PubMed]

Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception & Psychophysics, 63,* 1293–1313. [PubMed] [Article]

Wolfe, J. M., & Bennett, S. C. (1997). Preattentive object files: Shapeless bundles of basic features. *Vision Research, 37,* 25–43. [PubMed]

Xu, Y., & Singh, M. (2002). Early computation of part structure: Evidence from visual search. *Perception & Psychophysics, 64,* 1039–1054. [PubMed] [Article]

Yuille, A. L., Fang, F., Schrater, P., & Kersten, D. (2004). Human and ideal observers for detecting image curves. In S. Thrun, L. Saul, & B. Scholkopf (Eds.), *Advances in neural information processing systems* (Vol. 16, pp. 1459–1466). Cambridge, MA: MIT Press.