INVESTIGATION OF SHAPE REPRESENTATION USING SENSITIVITY TO AXIS

AND PART BASED TRANSFORMATIONS

by

KRISTINA DENISOVA

A dissertation submitted to the Graduate School-New Brunswick

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in

Psychology

Written under the direction of

Manish Singh

And approved by

____

____

____

____

New Brunswick, New Jersey

October, 2010

ABSTRACT OF THE DISSERTATION

Investigation of shape representation using sensitivity to axis and part based transformations

By KRISTINA DENISOVA

Dissertation Director:

Manish Singh

Part-based approaches organize global shape in terms of segmented parts and their spatial relationships, and are robust under transformations such as articulating limbs that are common in biological objects. It is well documented that transformations that alter qualitative part structure of a shape are more noticeable than those that do not. However, previous work has not compared sensitivity to transformations that change quantitative parameters involving the spatial relationships between parts, without altering the shape's qualitative part structure. Shape transformations were applied to a protruding part attached to a larger base (two-axial-branch shape: Experiment 1) or to a simple elongated shape (single-axis shape, Experiment 2): length, width, curvature, orientation, and location change of the part. Using a 2IFC task, increment thresholds were derived, and in Experiments 1 and 2 were converted into common metrics to enable comparison across transformations. Higher sensitivity was found for transformations involving the intrinsic parameters of a single axial branch such as length, width, curvature, and lower sensitivity for those involving two axial branches such as orientation, location, providing evidence for the single-part superiority effect. In Experiment 3, the orientation of the shapes (two-part and single-part) varied randomly on each trial. Results showed no benefit in sensitivity provided by presence of the base in the two-part shape in Experiment 1. In Experiment 4, the influence of surface, as opposed to contour, geometry was investigated by manipulating figure and ground regions defined by a fixed contour using stereoscopic depth. A protrusion (a positive part) in the figure region corresponded to an indentation in the complementary region (a negative part). Two conditions were tested, orientation and location; for each, sensitivity to positive and negative part changes was compared. Sensitivity was better for positive compared to negative parts in the case of orientation, indicating differential processing for transformations that are biologically plausible; no difference was found for location. Even when sensitivities were compared using a common scale, different transformations elicited very different sensitivities, indicating that purely physical measures of shape similarity are inadequate for perceptual similarity. The results also indicate that shape sensitivity depends on both contour and surface geometry.

Acknowledgement

I gratefully acknowledge the support and guidance of Manish Singh, my advisor and my committee members Eileen Kowler, Jacob Feldman, and Doug DeCarlo, for their insightful and helpful comments during the progress of this work, as well as Xiaotao Su for technical assistance. This research was supported in part by NSF IGERT Trainee grant.

Table of contents

List of tables

List of illustrations

1. Introduction

1.1. Vision: background

When we look around, whether trying to locate a suitcase on the conveyer belt at the airport, a mischievous kitten, or a friend in a busy train station, we are usually able to perform such tasks successfully. Our visual system has the capacity to perceive and locate objects and people in our environment even though what we are looking for may not be exactly the same as we have last seen it: for example, the suitcase arrived dented and misshapen but nonetheless recognizable. Thus vision can be said to perform many tasks crucial to an organism's survival such as to provide us with the necessary information regarding the objects of interest.

There are two overarching problems the visual system needs to solve. First, our perception of the world is three dimensional while the pattern of light intensities on the retina is two-dimensional. Second, the visual system must organize the retinal inputs into "perceptual units" such as objects and surfaces (they are not differentiated *a priori* at the level of the retina), and it must organize the representation of each individual object. The overarching problem for the visual system, then, is how to organize the input in terms of object surfaces and their shapes? This thesis is concerned with the perceptual organization and representation of shape. This is a critical function of vision since it provides information for the basis of interaction for living organisms: whether to approach an object of interest, such as food, avoid a predator, or predict how an object will look or behave in the near future.

1.2. Shape representation

How does the visual system represent the shapes of objects? In geometry, *shape* refers to those properties that are unaffected by rigid transformations (such as translation and rotation) or by uniform scaling. Thus, two shapes are geometrically "equivalent" if they can be brought into alignment by applying one or more of these transformations. Psychologically, this is not necessarily the case. First, geometrically equivalent shapes can look perceptually different. For example, when a square is tilted 45 degrees, it looks more like a "diamond" than a "tilted square" (Mach, 1914/1959; Figure 1a). Second, geometrically distinct shapes (not related by rigid transformations) can be perceived to be equivalent, such as various articulations of a human hand as illustrated in Figure 1b.

Theories of shape that do not involve a decomposition into parts are not able to effectively represent the kinds of global changes that normally occur in man-made and biological objects. For example, a template theory represents a shape in terms of a fixed pattern or template, and compares two shapes using point-to-point similarity *after* aligning the shapes as much as possible using rigid transformations (Ullman, 1989). While this theory is helpful for non-biological objects, especially in the context of mechanical applications such as comparing a template of a tool to a physical object (i.e., in looking, for example, for defects), it is not viable for biological objects, which undergo non-rigid transformations due to part articulation.

1.3. Part-based representation

In order to address these concerns, researchers have proposed that an object's shape is represented in terms of its component parts and their spatial relationships (Marr & Nishihara, 1978; Biederman, 1987; Hoffman & Richards, 1984). Phenomenologically, objects are perceived to be comprised of distinct parts. We effortlessly perceive the table in Figure 2 as an object with a top and four legs, a tree with its various branches, a frog with a trunk and four legs, and a kitten with a curled tail. From a computational perspective, the advantage of the part-based approach is that it separates the representation of the spatial relationships between the parts from the representation of the parts themselves. This separation makes part-based representations more robust with respect to the various articulating poses that the object may take: in one sense, the kitten may be said to have the same shape whether it is sleeping with its tail and limbs curled up, or running with outstretched limbs. This property addresses the main drawback of the template theory, namely, that shape's representation, including its parts' configuration, is fixed.

Given that the visual system's predisposition to divide shapes into parts, is there a systematic process by which it finds or segments shapes into parts? The shape primitives approach uses a fixed and predetermined repertoire of primitive shapes that can be used, in principle, to represent (or at least to approximate) any shape in the environment. Early versions of this approach use generalized cylinders (Binford, 1971; Marr, 1982; Marr & Nishihara; 1978) as shape primitives. A more recent version uses geons (Biederman; 1985,1987). This approach segments shapes into parts by fitting shape primitives to the image of an object. Unlike the template theory, any novel poses of articulated objects *can* be accommodated. The limitation of this approach however is that it is not flexible in a different way: the set of shape primitives is too limited to capture the rich variety of shapes we can see in the world.

A second approach to parts is based on using general-purpose geometric rules for segmenting shapes, such as the part-boundaries approach. One important assumption for this theory is that the visual system relies on world regularities. One such regularity is the transversality principle: when two smooth 3D objects intersect generically with one another, they produce a concave tangent discontinuity at their intersection as seen in Figure 3. Since these two objects now form a single object, if one wanted to individuate these two objects at a later time given only the composite shape, concave discontinuities would be good cues for finding separate parts. When concave creases are smoothed, they produce negative minima of curvature as illustrated in Figure 4. Based on these considerations, Hoffman & Richards (1984) proposed the minima rule, which states that one can use negative minima of curvature as boundaries for segmenting shapes into parts. This approach requires no familiarity with the objects (only geometric information is used). Moreover, it does not impose strong constraints on the shapes. That is, this approach does not require that shapes belong to any specific class (such as generalized cones). For example, in Figure 5, we see hills and valleys (valleys are along the dotted line; from Singh & Hoffman, 2001). However, when viewed upside-down, the hills and valleys reverse (dotted line now falls on the hills). The location (dotted lines) where the visual system segments the surface into parts changes because convex and concave sections of the shape switch, and part boundaries lie in concave (negative curvature) regions.

1.4. Evidence for parts in visual perception

There is extensive empirical evidence for part-based representation by the human visual system in various contexts. Influence of part structure on shape perception has been demonstrated, for example, in figure and ground reversal displays (e.g., Barenholtz & Feldman, 2006; Baylis & Driver, 1995; Hoffman & Singh, 1997; Stevens & Brookes, 1988), visual search (e.g., Xu & Singh, 2002), detection of shape changes (e.g., Barenholtz et al., 2003), and judgments of distance between points within shapes (Denisova et al., 2006). The computation of parts is fast, automatic, and thought to occur in parallel (e.g., Baylis & Driver, 1994; 1995; Hulleman, te Winkel & Boselie, 2000; Wolfe & Bennett, 1997). Additional evidence is provided by a series of single-cell recording studies. Pasupathy and colleagues (Pasupathy & Connor; 1999, 2001) found groups neurons in monkey's area V4 that specifically code for contour segments and are tuned to specific magnitudes of curvature, as well as sign of curvature (convex or concave), at different angular positions on the shape's outline.

Barenholtz & Feldman (2003) study investigated how the visual system segments shapes into parts using an objective perceptual judgment task and provided psychophysical support for Hoffman & Richards (1984) minima rule. In the experiment, the stimuli used were elongated shapes with varying degrees of curvature, such that the lowest level resulted in a unitary shape and the highest level segmented the shape into distinct bow-like parts (separate objects). Two marks were placed along the contour of the shape such that either a negative minimum or a positive maximum extremum was located between the two (the local geometry of both extrema was identical; see Figure 6). The subjects' task was to decide whether the two marks were the same or different and response time was measured. The results showed that within-part judgments were faster than between-part judgments for all curvature levels tested. This experiment controlled for the local geometry of the bounding contour of the shape and provided evidence that part boundaries are an important component of shape representation: the magnitude of the curvature was identical for curvature minima and maxima, and yet observers responded faster when the judgment involved a perceptual unit with boundaries at negative minima and slower when it did not (when the boundaries lay at positive maxima).

Cohen & Singh (2006) investigated the effect of manipulating a part's salience (i.e., the sharpness of its boundaries) on visual system's judgments of overall orientation of a two-part shape (a shape with a part protruding from its base). The global orientation of a shape with part structure can be computed by the visual system in accordance with two hypotheses: either all points within the shape are treated uniformly, regardless of whether the shape has part structure (homogenous computation: the global principal axis is computed), or the parts may be assigned different weights (part-based computation: differentially-weighted principal-axis computation). The shapes were constructed such that the part's salience depended on curvature at the boundary between the attached part and the base (higher curvature denoted higher part salience or independence); different sizes of the parts (small to large) were also tested. Observers participated in an adjustment task (Experiment 1), and were instructed to adjust the probe line to match apparent orientation of the whole shape (see Figure 7). Results showed that observers' settings approached the base-part axis with increase curvature at the part boundaries—as the part becomes more distinct from the base. When the curvature is minimal, and the part is virtually indiscernible from the base, the adjustment settings are close to the entire shape's principal axis. Similar results were obtained in Experiment 2 which used a 2AFC task: perceived orientation was closer to the principal axis of the base part for sharp curvature at the part boundaries, and closer to the global principal axis with smaller turning angles. This experiment reveals a systematic pattern: when the part of the shape is highly salient, observers disregard the attached part, and judge the overall orientation with respect to the base part only. However, when the part is not sharply delineated from the base (weak boundaries), observers judge orientation according to the shape's global principal axis. This is consistent with a differentially-weighted principal-axis computation in which the attached part is assigned a lower weight with increasing part salience.

De Winter & Wagemans (2006) investigated how observers segmented 2-D outlines of common objects into parts. Their stimuli were a set of outlines of real-world objects (a total of 88 outline shape; 22 per each subject). Observers were instructed to first try to identify the outline (what object it represents), and then to draw segmentation lines so that they represent important or salient parts. The segmentation lines could be either straight or curved. Results showed that negative minima were the most common segmentation points; inflections and positive maxima were less common. This pattern was the same for both easy and difficult segmentation conditions (Figure 8 shows a sample object, a wineglass). It is necessary to note however that subjects' segmentation cuts in this task are a combination of both higher-level—cognitive—and low-level processes.

In the context of a figure and ground decision task, Barenholtz & Feldman (2006) investigated the visual system's expectations about an object's possible deformations based on its geometry. The researchers proposed that, in ambiguous figure-ground contexts, there is a tendency to assign figure such that articulations are perceived as occurring at concavities. In Experiment 1, observers were presented with a centrally located ovoid shape divided into alternating color regions, as shown in Figure 9a (the researchers used a stationary contour point as a hinge between rotating edges). The task was to indicate which color appeared to move; motion was achieved by rotating one of top edges through 10 angular degrees to the left or to the right about the vertex. In the experiment, the sharpness of the vertex angle was manipulated and it was predicted that a more acute angle (i.e., a sharper part boundary) would result in greater tendency to assign "figure" to that color region. Results indicated that as the angle became sharper and more concave observers exhibited a stronger bias to assign figure such that its articulating vertex is concave. Barenholtz & Feldman (2006) performed a series of experiments in which the strength of this dynamic articulation cue was tested relative to other known figural cues, such as the width of the color regions (narrow or wide). In Experiment 2, observers showed consistent bias in assigning narrower stripes as 'figure', compared to wider stripes (static condition; Figure 9b). However, when the dynamic articulation cue was also used (motion condition), the overall bias to assign narrower stripes was reduced, and even reversed in some cases. Observers preferred an articulating-concavity figure assignment, even if this led them to interpret a wide stripe as figure. Results showed that subjects' judgments indicate bias towards assigning figure so that motion is consistent with convex parts articulating at concave minima, and this bias can sometimes override well-known Gestalt biases for assigning figure and ground.

Cohen & Singh (2007) used an objective segment-verification task to investigate the way in which an object is segmented into parts that are defined by geometric attributes. In the experiment, subjects were presented a probe shape, followed by a mask, and then a test segment, and asked to indicate whether the test segment matched a portion of the probe shape (probe and test shape are depicted in Figure 10). In this study, the critical manipulation was that the segments were bounded by either negative minima or positive maxima of curvature. Since the segments produced by both types of boundaries constitute legitimate subsets of the shape, the goal of the experiment was to find out whether accuracy is greater for either type of the segment. Specifically, the prediction was that the segments that constitute “natural” parts or “units” of the shape will result in better performance on the verification task. The results showed that accuracy was higher for those segments whose boundaries were defined by negative minima rather than positive maxima. Follow-up experiments confirmed that this result is not due to segments’ variations in length between the maxima- and minima-bounded segments (Experiment 1b), and the inflections were not found to play a significant role in part segmentation (Experiment 2). Thus observers in this study were asked to verify whether a given contour segment is a component of the shape (without explicitly requiring observers to indicate part segmentation). The results indicated that the visual system is more readily able to identify whether the segment is a part of the shape if it is bounded by negative minima.

1.5. Axis-based representation

A skeleton or axis-based approach uses a compact stick-figure representation of a shape to capture its qualitative geometry, e.g., the branching structure of its parts. It is an efficient representation that is able to capture the basic structure of the shape without necessarily representing details of contour geometry. For example, the "pipe-cleaner" object in Figure 11 lacks surface texture, color, and additional properties, and yet it is easily recognizable as a giraffe (from Marr & Nishihara, 1978). In other words, such representations make the essential internal structure of the shape explicit (Blum, 1973; Marr & Nishihara, 1978; Leyton, 1989; Kovacs et al., 1998). Furthermore, this approach can allow one to establish a one-to-one correspondence between axial branches and parts (Feldman & Singh, 2006). In Computer Science, many algorithms for computing part structure are also based on axial representations (e.g., August, Siddiqi & Zucker, 1999; Mi & DeCarlo, 2007; Mi, Stone & DeCarlo, 2009; Siddiqi & Kimia, 1995).

Blum (1973) introduced the Medial-axis transform (MAT) as a way of extracting an object's symmetric axis. The medial-axis of an object is extracted using the locus of the centers of maximal disks that can be inscribed within the shape as depicted in Figure 12. This axial structure turns out to be equivalent to using a "grassfire" process. Given a particular shape such as a rectangle, a grassfire process starts at the edges of a shape all at once, and burn towards the center; this process results in internal skeletal axes. However, MAT has been critiqued that it is neurobiologically intensive (and thus not neurologically plausible), and highly sensitive to variations in the figure's contour (computation of axis is not robust to random perturbations; e.g., Kovacs et al., 1998; Feldman & Singh, 2006).

Investigations into the role of axes in visual representation of shape have included both informal examinations and empirical studies. For instance, Psotka (1978) asked observers to place a dot inside an outline of a shape at the "first location that comes to mind". Each observer placed a single dot on each shape, and data from one hundred and fifty people (Experiment 1) were superimposed on the experimental shape outlines. The patterns of the dots falling within the shapes corresponded to predictions given by the Blum's grassfire model: For example, a rectangle produced a pattern of a central line flanked by two outward lines, and a human figure shape resulted in pattern corresponding to the skeletal bones of the trunk and the arms. Note that for the human figure, the dots were not placed uniformly along the skeletal structures, but were also guided by cognitive expectations about importance of internal structures: for example, more dots were placed in the area corresponding to the heart, and fewer near the feet. This exploratory experiment shows that observers have the capacity to make judgments about shape in accordance with Blum's hypothesis emphasizing axial structure, however, more recent rigorous studies provide support for axial representation that are less prone to cognitive factors.

Burbeck and Pizer (1995) focused on identifying important regions at different spatial scales in an image using their *core* analysis. The model rests on the idea that first, borders of regions (outline contours) are signaled by "boundariness" detectors: the neurons that indicate the presence of a boundary and that are sensitive to orientation and scale. The boundariness detectors connect with one another across the shape in the middle of a particular region, providing a "medialness" measure at a particular scale. For example, for a tear-drop object in Figure 13, the location of the core is the figure's "mid-line"; the width of the core is represented by the "fuzz" around the core with wider regions exhibiting greater amount of fuzz (greater degree of imprecision in the representation of its location). The location of the core is represented by the associated boundariness detectors at a particular spatial scale. Therefore, two different objects that are similar on a large scale would be represented by similar cores (when boundariness detectors on a large spatial scale are considered); however, if they differ on a small spatial scale, such as when one's contours are jagged, boundariness detectors active at these spatial scales will produce smaller cores, providing a way to represent the difference in these two objects' bounding contours. The core model's strength lies in the fact that minor perturbations along the object's edge would not affect the core that represents the coarse spatial scale analysis of the object. However, Burbeck and Pizer (1995) note that the core model has certain weaknesses. For example, the core represents the object's middle and is related to that object's major axis, but the process of deriving the core rests on the ability to locate the boundary, which may not be available in all cases due to occlusion of the edge.

Kovacs and colleagues investigated shape representation theoretically and empirically using the medial-point analysis (e.g., Kovacs & Julez, 1994; Kovacs, Feher, & Julez, 1998). Their procedure first finds the medial axis of the shape, and then isolates the most "informative" point along that medial skeleton, which they argue are the branching points between axial branches. The authors note that this method aims to capture a description of shape that is local, compact and is available across different spatial scales. Their experimental displays consisted of a field of elements (Gabor patches) that were all defined by the same parameters except one, the target, that had a different contrast relative to the background elements. The background elements were oriented randomly, with the exception of a closed contour embedded within the display. The contour shapes were varied and included a circle, an ellipse (with varying aspect ratios), cardioid-shaped and triangle-shaped contours. In a 2AFC task, observers were asked to discriminate a centrally located target that differed from the background elements in its contrast. The location of the target relative to the embedded contour varied across experimental trials. Results indicated that observers' sensitivity to detect the target was affected by the target's location within the region enclosed by the embedded contour. Specifically, for the ellipses, there were two enhancement regions that elicited greater sensitivity and were located symmetrically from the perimeter of the embedded contour. The authors note that there is a high correlation between the peak locations predicted by the model, and the measured locations of peak sensitivity. This approach to shape representation is compact, since a small number of points (1 for circle, cardioid, and triangle shape, and 2 points for an ellipse) is required. In effect, this approach takes the Medial axis as the starting point and locates points that are particularly informative.

Axial-based representations are able to capture representations of objects, including biological ones in a compact and efficient way, and carry information about their essential qualitative structure. In addition, such representation can reveal the paths along which forces acted in generating the shape (Leyton, 1988). Recent computational approaches (e.g., Feldman & Singh, 2006) improve some deficiencies of the axis-based approach (e.g., increased robustness to noise, and establishing a one-one correspondence between axial branches and parts).

1.6. Sensitivity to changes and its implications for studying shape

A fruitful way of studying shape is by measuring the visual system's sensitivity to change. In change detection studies observers are asked to judge whether sequentially presented images are similar or different, and measure what types of changes are more or less easily detectable. Detection of change has been studied in different contexts. For example, previous studies investigating change detection in complex scenes found that observers are often unable to detect substantial changes within such scenes ("change blindness"; e.g., Rensink, O'Regan, & Clark, 1997).

In contrast, more recent studies have investigated whether certain aspects of shape are represented differently (e.g., emphasized to a greater extent) than others by the visual system. It has been noted in the literature that, for reasons of parsimony, certain aspects of the scene are represented more explicitly by the visual system, while others are relegated to the background, or represented only implicitly. For example, when considering a shape's contour, not all points along the contour may be equally informative (Alhazen, ca. AD 1030; in Sabra, 1989). Attneave (1954) argued that points of maximum curvature are most informative for recognizing shape. When Atteneave (1954) asked subjects to mark the most salient points along a shape's contour, most marked points were positioned around points of curvature extrema—those with locally highest magnitude of curvature. Not all studies empirically supported this outcome, however: a study by Kennedy & Domander (1985) found that identification of stimuli was most effective when fragments of stimuli were placed midway between points of extrema. Yet additional follow-up studies seem to confirm Atteneave's (1954) hypothesis. For example, Norman, Phillips, and Ross (2001) asked observers to mark ten points along a shape's contour that they deemed most "salient", meaning that the new contour consisting just of these points resembled the original version most closely. The authors found that the points along the contour found to be most informative by the observers correlated with points of curvature maxima. Furthermore, not all curvature maxima may be equally informative, however, as several studies indicate. For instance, Feldman & Singh (2005) note that negative curvature extrema (minima) are more informative than positive extrema because closed contours are necessarily more convex than concave.

A series of empirical change-detection studies (e.g., Barenholtz et al., 2003; Cohen et al., 2005) used a shape (a randomly generated polygon) that changed in a particular way on any trial, and asked observers to judge whether the shape had undergone a change between the two successive presentations. In Barenholtz et al. (2003) study, the base shape was modified by either adding or subtracting a vertex (Figure 14). The vertex that was added or removed could either be a convex vertex or a concave vertex. Observers were presented with two shapes: an original and a potentially changed version (50% of the trials were "changed" trials; the "non-change" trials contained identical shapes for both display intervals). Results showed that observers were more accurate (accuracy expressed in terms of *d*') to detect change when the shape changes involved a shape modification to a concavity rather than a convexity; this effect increased with the magnitude of the vertex change.

The differential results of Barenholtz et al. (2003) could be due to the fact that changes involving addition of concavities also changed overall object's part structure in a more significant way than did convexities; for example, a large magnitude of change at a concave vertex can affect the number of perceived parts (i.e., it can introduce a new part), while introducing a convex vertex (as well as removing a convex vertex) does not affect the number of perceived parts. Alternatively, the difference in sensitivity could truly be due to differential underlying representation of local convexities and concavities themselves.

In order to investigate this possible difference in representation further, Cohen et al. (2005) also used a detection task to gauge possible differential sensitivity of observers to shape changes. Cohen et al. (2005) investigated this issue under conditions where the overall, gross structure of the shape was constrained while an existing convex or concave vertex (i.e., not introducing a new one) was either enhanced or diminished. Thus, as can be seen in Figure 15, a given shape change did not fundamentally affect the part structure of the entire shape. The task was to judge whether there was a change in shape between two subsequent presentations of the shape. Results indicated that observes were significantly more likely to respond correctly when concave, rather than convex, changes were involved. Thus observers were more sensitive to concave compared to convex types of changes, even though neither type of change qualitatively affected the overall part structure of the shape.

A follow-up experiment (Experiment 2) controlled for geometric differences between the stimuli. This was accomplished by creating two shapes that shared the same contour (e.g., Atteneave's divided egg; Figure 16). Each half had identical contour, but opposing sign of curvature, and each contour is presented twice, once with each of the figural sides of the ellipse, thereby equating geometric differences in the stimuli. In addition, the "critical" vertex and the neighboring vertices of the vertex had the same sign of curvature (if the vertex underwent concave change, then flanking vertices also were of concave type, and vice versa). Thus a critical vertex was convex (and nearby vertices convex), and a complementary critical vertex was concave. This manipulation ensured that the flanking vertices were not, for example, concave, when a convex change was applied; the changed vertex was thus always an "apex of a pentagonal sequence of vertices" (Cohen et al., 2005). This experiment supported the findings of Experiment 1: detection accuracy was significantly higher for concave than convex vertices.

A recent study by Barenholtz & Tarr (2008) investigated similarity judgments to different kinds of shape transformations. Observers were shown three shapes: a standard, and two transformed versions of the standard. One of the transformations was a 'biologically valid' articulation of the standard object shape's part, and the other 'biologically invalid'. A total of two kinds of transformations were used. In the *figure/ground* reversal condition, two different shapes were created that shared the same contour. Whether articulation of the part is considered 'biologically valid' depends on whether the part rotates (articulates) at concavities, when it is presented as figure, or convexities (when contour polarity reverses and figure becomes 'ground'; shown in Figure 17). In the *axis-of-rotation* reversal condition, the shapes differed by the location of the part's fulcrum. In the 'biologically' valid version, the part articulates at negative minima (this arrangement preserves the part's location with respect to the base shape). In the 'biologically invalid' version, the fulcrum of the part's articulation is at the end-point of the part. Such articulation preserves the part's orientation overall with respect to the base shape (it is identical as that in the 'biologically valid' version), but it does not preserve the part's boundaries, where part meets the base; this idea is illustrated in Figure 18.

Results showed that subjects preferred a "valid" articulation, for both types of reversals, on 70% of the trials, indicating that shape transformations that are biologically valid are viewed as more similar (to the standard shape), versus a geometrically comparable, but biologically invalid transformations. This outcome indicates that subjects preferred those transformations that preserved part-shape location—the spatial configuration of the part's attachment to the base is constant. The authors note that this outcome is in line with Hoffman & Richards (1984) minima rule and the notion that it may be helpful to divide an object into its constituent parts since the spatial relationship between the part and the base (in this context, specifically, the location of attachment) is constant across biologically valid articulations, but is not preserved when articulation changes the location of attachment of the part to the base ("biologically invalid" articulations). Moreover, this study supports the idea that the visual system encodes explicit relational and configural information about a shape, and not simply some featural properties about the shape, or configural information along with featural properties (e.g., Zhang & Cottrell, 2005).

In summary, previous studies found that sensitivity to detect changes in shapes is higher when concave vertices are involved in the changes (e.g., Barenholtz et al., 2003). However, a series of studies by Bertamini and colleagues also used change detection methodology, and found a somewhat different, context-dependent outcome for detection of different types of vertices. Bertamini & Farrant (2005) investigated whether there is higher sensitivity for changes to either concave or convex vertices depending on the context in which they are presented (e.g., higher only in the case where a concave vertex may affect part structure of the shape, and thus facilitates detection—a contextual effect). In Experiment 1, observers were shown two polygonal shapes. The shapes could be presented as figure or as a hole within a square. For each type of stimulus, a contour change was introduced for either the first or the second of the two shapes: a new vertex was introduced which could be turned either inwards or outwards; this new vertex replaced a straight segment in the shape. When the shape was shown as figure, the inward-turning of the vertex was concave, and the outward-turning was convex. This relationship reversed in the case of the hole: the inward-turning vertex became convex, and outward-turning—concave. These manipulations are illustrated in Figure 19. The shapes were displayed using random dot stereograms, ensuring figure-ground reversal but keeping the contour congruent. This allows one to disambiguate the context within which the shapes are presented; performance is predicted to be higher for the concave (inward turning) vertices in figures and for convex (outward turning) for holes. One of the shapes was a changed version of the other shape, and the observers' task was to indicate whether the shapes were the same or different. Overall, performance was higher for figures than for holes, and sensitivity was higher for the "turning inward" (i.e., concave) condition. The trend was present for both figures and holes, and statistically significant for the figures (i.e., higher sensitivity for concave vertices), but not for holes. For the holes, the opposite trend was expected (higher sensitivity for convex vertices). Experiment 1 indicates that in contrast to previous work, sensitivity to concavities is not higher in *all* contexts, since no effect was found in the case of holes. Bertamini & Farrant (2005) note that one issue is that there are more convex than concave vertices in the polygon figures used in this experiment (this is the case for all closed shapes: such shape may have only convex vertices but cannot have only concave vertices). Thus, adding a concave vertex may appear more salient than a convex vertex. In the context of holes, there are more concave vertices than concave and theoretically adding a convex vertex may be more salient than a concave one (although the results of Experiment 1 did not confirm the prediction for the holes case). Experiment 2 addressed this problem of isolating a part of the whole object by displaying it through an aperture. In this way only local information is available to the visual system. To do this, two types of *contexts* and two kinds of vertices changes were tested. A *convex context* consists of a "hill" area topped by a straight segment. This segment can either protrude away from the hill, constituting a convex change (A; "a convexity changes shape") or into the hill (B, "a part splits into two"), a concave change. A *concave context* consists of a valley, with a straight segment on its floor. This segment can also either protrude outwards (C; "A new part is introduced"), a convex change, or inwards (D; "A concavity changes shape"), a concave change. Different predictions are made for these cases: 1) higher sensitivities for the cases for B and C, when the sequence of concavities and convexities changes (i.e., the "bracketing hypothesis"); the authors refer to "bracketing hypothesis" as the idea that part segmentation is likely to occur when convexities and concavities "follow each other" on a contour (i.e., part segmentation does not necessarily occur at negative minima). The alternative prediction, 2), is that higher sensitivities are expected for B and D if minima points are salient in general (consistent with previous work). The stimuli used in the experiment constituted a portion of a larger shape with a fixed number of vertices occluded by a frame on three sides; the figural side of the contour could appear either on the left or the right.

Results showed that sensitivity was higher for changes between convex vertices (A&B), compared to C&D. There was also a significant interaction between type of vertex change and context: performance was higher when there is a change of curvature sign. Specifically, sensitivity was higher for the B case (a new concave vertex) but only within the Convex context, while sensitivity was higher for the C case (a new convex vertex) but within the Concave context. Note also that additional post-hoc tests did not reveal statistically significant difference between B and C, and between A and D, indicating an absence of a benefit of either convexity or concavity vertices when context is taken into account. On the basis of this study, Bertamini & Farrant (2005) suggest that concavities may be salient under some conditions, but not others, for example, since new convex vertices are more salient in the concave context condition. This is predicted by the bracketing hypothesis, which dictates that concavities are more salient only when they affect part structure.

Another change detection study investigated the question of how objects are segregated from the background: which regions tend to be assigned figural status? The notions of convexity and concavity of a contour are defined with respect to whether the contour belongs to the figure or the background. There is both evidence that the assignment of figure status may be biased to the convex side (Kanizsa & Gerbino, 1976) as well as evidence that concavities are more salient than convexities: e.g., Barenholtz et al. (2003). Bertamini & Lawson (2008) used random-dot stereogram displays which ensured unambiguous specification of which region appears in front (defined by binocular disparity). The regions had a convex or concave contour, only a portion of which was visible to the observer through a square aperture (this method effectively isolates the critical contour). Faster responses are predicted for convex regions. This is so because there is a bias to perceive convex regions as figures, and figures tend to be perceived in front of the background. If the binocular disparity is consistent with this expectation (e.g., convex as figure, and figure actually in front), then the responses are expected to be faster. The responses are thus expected to be slower in the case when the convex region is in the back, since the binocular disparity given to the convex region conflicts with the expected figural (and thus in front) assignment (e.g., convex, but as background). In Experiment 1, observers' task was to judge whether the right or the left surface was in front. Experimental manipulations are illustrated in Figure 20. Results showed that responses were faster when the convex region was in front, as well as when the larger disparity was used between figure and ground regions. This confirms the prediction: when convex region is in front, responses are faster compared to when concave region is in front. Experiment 2 was performed to address the possibility that when the convex region is in front, and observers look at the center of the display, observers fixate the front surface of the convex region. However, in the case of concave regions, observers would fixate the ground (i.e., centre of display) when the concave region is in front. This confound occurs since the two regions are matched for area, and may lead to faster responses for the convex regions. Thus, the second experiment introduced several additional conditions: first, stimuli with straight contours that divided the area into smaller and larger portion (either the larger or smaller portion could be specified in front). If the extent of surface area is taken into account, then those regions with greater surface area should yield faster responses. Second, the horizontal extent of the stimuli with convex and concave contours was varied. The stimuli with greater extent of the curved portion should also yield faster responses because observers would be looking at the front surface even more for this type of stimulus. Results showed that there was no difference between the straight-contour stimuli: i.e., if a larger region was in front it did not provide an advantage over when the smaller region was in front. As in Experiment 1, responses were faster for convex-in-front stimuli. However, no effect was found for the amount of horizontal extent of the curved stimuli. This outcome may be related to Feldman & Singh (2005) notion of "surprisal value": the information content is greater at concavities because concavities, on average, are less likely to occur than convexities.

1.7. Sensitivity and Similarity

The differential sensitivity to various shape transformations reported in previous studies on change detection also has clear implications for similarity between two shapes. If observers are very sensitive to a particular type of change, then a small amount of shape change will be sufficient for the shapes to look dissimilar. On the other hand, a lack of sensitivity to a particular shape transformation will cause two shapes to continue to look very similar even with a large physical change. As noted above, previous work has shown that changes that alter qualitative (i.e., skeletal or part) structure of the shape lead to shapes that look dissimilar from one another. For example, Figure 21 shows that different transformations of a base shape (an elongated tube) that involve a comparable amount of physical change can be perceived as more or less similar to the original shape. When a small bump is added to the elongated tube, the result is a perceptually very different shape. However, another transformation which merely elongates the tube has little perceptual effect: that shape looks similar to the base shape. However, even though the elongation transformation is perceptually less striking than the "bump" transformation, the elongation change is physically more different in terms of an objective metric such as the two shapes' non-overlapping areas.

How can the psychological similarity between shapes be measured? Previous studies investigated the degree of similarity between stimuli using a variety of methods, such as ratings or rankings. For example, the multidimensional scaling (MDS) technique (from Shepard 1962a, 1962b) constructs representations across one or more relevant dimensions in a pair of objects. The similarity is measured in terms of distance between the corresponding points in a multidimensional space. If one were to measure similarity between several colors, the colors would be arranged in a proximity matrix (Palmer, 2001, p. 388) and observers would be asked to give a rating about the extent to which any two colors appear more or less similar to one another (e.g., violet is likely to be rated highly similar to blue, but highly dissimilar to yellow). The output is related to these similarity ratings: the greater the rating between any two dimensions, the closer they should be arranged in the multidimensional space. Psychophysical studies have measured shape similarity by asking subjects to the rate the similarity between pairs of shapes (e.g., Briscoe, 2008), and tested sensitivity to different shape changes using change detection methods (e.g., Barenholtz et al., 2003; Cohen et al., 2005).

Another perhaps more sensitive psychophysical method to measure observers' sensitivity to shape changes may be to obtain difference thresholds. The threshold would indicate how much difference between two shapes is necessary for observers to detect a difference between the two shapes. Thresholds have been obtained in the context of relatively simple parameters of stimuli. Previous studies investigated observers' sensitivity to line length, curvature, and orientation in the context of simple line segments. It has been shown that observers are able to discriminate a 0.15 to 0.5 degree difference in the orientation of two successively presented lines (e.g., Burbeck & Regan, 1983). Observers are also good at judging the difference in angles formed by the two intersecting lines (the "Vee" angle), resulting in just noticeable difference ranging from

0.7 to 1.6 degrees (Regan & Hamstra, 1992). The accuracy of judgments for these stimuli was better for 90 degree angles than for other arbitrary angles (Gray & Regan, 1996). Furthermore, Watt & Andrews (1982) investigated curvature discrimination for curved line segments and found that thresholds corresponded to maximal differences in retinal position of less than 5 seconds of arc. The authors emphasized the importance of disambiguating effects of line length and orientation change on curvature discrimination (for example, when line length was increased, threshold for curvature also increased).

As noted above, previous work has shown that transformations that alter part structure (i.e., the shape's qualitative branching structure) are more easily detected than those that do not. However, previous work has not systematically investigated shape transformations that preserve qualitative axial/part structure, while manipulating specific quantitative parameters in an axial representation. The current experiments will be applying shape transformations that involve manipulating quantitative variables in a shape's axial representation (e.g., part length, width, orientation, curvature, and location) that do not alter the qualitative axial or part structure, and comparing sensitivities to these different transformations. Specifically, this study will obtain difference thresholds for changes along various axis-based parameters that quantitatively measure the perceived difference between the two shapes.

The pattern of differential sensitivity will facilitate development of measures of shape similarity that are perceptually informed. For example, information about how quantitative changes to various shape parameters affect perceived difference between shapes is important because psychological similarity can be different even though objective, physical differences are the same. This information will thus be relevant to computer vision researchers who are also interested in developing measures of shape similarity.

2. Experiment 1: Two-part shapes

The main goal of this Experiment is to measure and quantify the differences in sensitivity of human observers to several types of transformations in a shape with part structure. As noted earlier, axis-based descriptions of shapes are capable of robustly representing the shape of various objects in the environment, both inanimate, and biological ones (e.g., Blum, 1973, Feldman & Singh, 2006; Kovacs & Julez, 1994; Kovacs, Feher, & Julez, 1998; Leyton, 1988; 1989).

The strategy used in this experiment is similar to the change detection paradigm. This methodology has been used to study the psychological similarity between two shapes (e.g., Barenholtz et al., 2003; Bertamini & Farrant, 2005; Bertamini & Lawson, 2008; Cohen et al., 2005). For example, if a shape needs to be changed a great deal before the difference between the transformed shape and the original version can be noticed, it suggests that the visual system is not particularly sensitive to that type of transformation. On the other hand, if the transformed version is changed by a small amount, and yet it is perceived as noticeably different from the original version, it suggests that the visual system is more sensitive to that transformation. However, in contrast to the previous studies on change detection on shapes, this experiment measures increment thresholds for different types of shape transformations (along the lines of psychophysical research that measured sensitivity to changes in simple stimuli such as line segments; e.g., Burbeck & Regan, 1983).

A total of five transformations were tested on a two-part shape composed of a part attached to a larger base, with each shape transformation applied only to one part, or one axial branch, of the two-part shape. Specifically, the transformations of length, width, curvature, and orientation were applied to the part of the shape, as well as a transformation of lateral shift in the location where the part protruded from the base shape. The goal was to compare perceptual sensitivity to these transformations. On each trial, observers saw three shapes on the computer monitor: a "test" shape followed by two alternative shapes. One of the alternatives was a transformed version of the shape, while the other was exactly the same as the test shape. Using a 2IFC task, observers were asked to make a perceptual judgment: which of the two alternatives matched the test shape—first or second?

2.1. Method

2.1.1. Observers

Six graduate students from Rutgers University participated in the study. Five of the observers were naïve about the experimental goals and were paid volunteers; one was the author. All had normal or corrected-to-normal vision.

2.1.2. Apparatus

Stimuli were generated using MATLAB (Mathworks) and the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997). Stimuli were presented on a high resolution (1280 x 960) 19-inch monitor (Mitsubishi DiamondPro) with a 120 Hz refresh rate, connected to a dual-core G5 Macintosh computer.

2.1.3. Stimuli

The shape used in this experiment was composed of an upward-facing protruding part, which was attached to a larger horizontally-aligned base shape. When the axis of the attached part was straight, it was a rectangle with smoothed corners (using a 1-D Gaussian smoothing operation). The part's width subtended .86 degrees of visual angle (dva) and its length subtended 2.58 dva. The base of the shape was a larger rectangular shape with rounded corners, with its length subtending 8.6 dva and its width, 1.2 dva. An example of the test shape is shown in Figure 22 (1^{st} column, 1^{st} row); a 1-D Gaussian smoothing operation was applied to the 4 corners of both part and base shape as well as to the junctions (part boundaries) between the attached part and the base shape.

The shape was centrally positioned on the monitor. The part was attached to the base approximately in the middle of the base. The exact position of the part varied randomly within the range of +/- 1.55 dva away from the vertical meridian on any given trial (this range is approximately equal to the 1/3 of the length of the base shape) for all conditions, except the lateral shift (location transformation) condition. In this condition the placement of the attached part had one of 2 values, either +1.20 or –1.20 dva away from the center of the base shape. In addition, the attached part was oriented either + 20 or – 20^{o} away from the vertical meridian in order to avoid judgments strictly at vertical meridian (in all conditions except the curvature condition).

There were three versions of this shape: a "test" shape and two "alternative" shapes. The parameters of the test shape for length, width, orientation, and location transformations were exactly as described above (these transformations all involved shapes with straight axes). In the curvature condition, the test shape was different in two respects: first, it had a baseline orientation of 0^o (it was vertical at its lower end where it attached to the base shape; this shape is illustrated in the 1st column, 4th row: Figure 22). Second, its baseline curvature parameter was set to .41 (1/dva) (the axis of the shape was an arc of a circle with the radius equal to 1/ 0.41, i.e., 2.44 dva and could curve either to the right or to the left of the vertical meridian of the monitor). One of the two alternatives resulted from applying a particular transformation of the test shape, while the other alternative other was identical to the test shape. The part transformations tested were length of the part, its width, orientation, curvature, and lateral shift (i.e., location of attachment of the part to the base shape).

In the length condition, the alternative shape was different from the test shape in the extent of its elongation. The length of the shape was manipulated by applying one of seven increments to the initial length parameter of 2.58 dva: .09, .17, .27, .34, .43, .52, & .60 dva. The length transformation is illustrated in the first row of Figure 22.

In the width condition, the length of the ribs (i.e., width) of the part of the alternative shape was modified; seven increments were applied to the initial width parameter of .86 dva: .03, .05, .09, .12, .15, .19, & .22 dva. The width variable corresponds to approximately twice the rib length in an axis-based representation. The width transformation is illustrated in the second row of Figure 22.

The orientation transformation modified the part's tilt away from the baseline orientation; seven increments were applied to the baseline orientation parameter (i.e., +20 or - 20^{o} from the vertical): 2, 6, 10, 14, 18, 22, & 26 degrees. The orientation transformation is illustrated in the third row of Figure 22.

The curvature transformation changed the extent of the bending of the part's axis. One of seven increments was applied to the base curvature parameter of 0.41 (1/dva): .06, .12, .17, .23, .29, .35, .41 1/dva (for subjects KD, SK, & CC) and .04, .08, .12, .16, .2, .24, .28 1/dva (for subjects SC, SS, SHK). These different ranges were selected for different observers based on pilot experiments. The curvature transformation is illustrated in the fourth row of Figure 22.

The location condition shifted the location of the part relative to the base shape. One of seven increments was applied to the initial location of the part relative to the horizontal center of the shape (either +1.20 or –1.20 dva): .14, .28, .42, .55, .69, .83, or .96 dva. The location transformation is illustrated in the fifth row of Figure 22.

2.1.4. Procedure

Observers viewed the displays binocularly with the head position fixed by a chinrest. Observers were instructed to look at the fixation cross which was presented at the center of the monitor at the beginning of each trial. As shown in the trial sequence in Figure 23, a total of three shapes were presented on each trial: the test shape and two alternative shapes. Each shape was shown for 200 ms and was followed by a mask in order to preclude an afterimage. The test shape and its mask were followed by a longer 900 ms pause; the two alternative shapes were separated by 300 ms. The final frame of each trial was a mask display, which remained on the screen until the observer made a response on the keyboard. One of the alternative shapes was identical to the test shape, while the other alternative was a transformed version of the test shape. The order of these alternatives (i.e., either the identical shape first, followed by the transformed alternative, or vice versa) was determined randomly.

The 2IFC task of the observer was to indicate which interval, first or second, contained the shape that matched the test shape. The observers were instructed to press '1' on the computer keyboard if the first alternative matched the test shape, and to press '2' if the second alternative matched the test shape. The observers' responses were not timed, and no feedback was provided after the response. After the response was given, the fixation cross reappeared to signal the onset of the next trial. Observers had the option to terminate the experimental session at any time by pressing the "escape" key. None of the observers terminated any of the experimental sessions.

2.1.5. Design

Each observer participated in a total of 10 experimental sessions, 2 sessions each for each of the five transformations of length, width, orientation, curvature, and location (within-subjects design). Each transformation condition contained a total of 350 trials divided into two sessions (one session contained 140 trials and the other 210 trials). This resulted in a total of 50 repetitions for each of the seven increment values for each transformation. The order of the conditions was counterbalanced across observers. In addition, each observer completed one practice session prior to each transformation session (140 trials).

2.2. Results

Weibull psychometric curves were fit to individual observers' data for each transformation condition using the psignifit software for MATLAB (Wichmann & Hill, 2001a). Based on these fits, increment thresholds and corresponding 95% confidence intervals were computed (i.e., increment needed to respond at 75% accuracy). The raw data (proportion correct for each increment level) and the psychometric fits are shown in Figure 24a-e, for each condition. Most observers' responses range from chance at the lowest increment value for each condition, to around 95% for the highest. Table 1 provides an overview of the thresholds.

<u>Increment thresholds</u>. Thresholds for length ranged between .21 [.17 .24] dva (observer SC) and .39 [.34 .42] (CC). Other observers' thresholds were .22 [.16 .25] dva (SS), .23 [.19 .26] dva (SHK), .30 [.28 .36] dva (SK), and .30 [.25 .33] dva (KD).

Thresholds for width ranged between .05 [.04 .06] dva (observer SHK) and .11 [.10 .13] dva (KD). Other observers' thresholds were .06 [.05 .07] dva (SK), .06 [.05 .07] dva (SC), .07 [.06 .08] dva (SS), and .09 [.08 .11] dva (CC).

Thresholds for orientation ranged between 6.08 [4.89 6.96] degrees (observer SS) and 10.7 [9.47 12.00] degrees (KD). Other observers' thresholds were 6.76 [5.61 7.59] degrees (SK), 9.45 [8.19 10.42] degrees (SC), 10.14 [8.67 11.16] degrees (CC), and 10.43 [9.23 11.22] degrees (SHK).

Thresholds for curvature ranged between .09 [.08 .10] 1/dva (SC) and .16 [15 .18] 1/dva (KD). Other observers' thresholds were .09 [.08 .11] 1/dva (SHK), .09 [.08 .12] 1/dva (SK), .10 [.08 .12] 1/dva (CC), and .12 [.10 .13] 1/dva (SS).

Thresholds for location ranged between .24 [.13 .27] dva (observer SK) and .62 [.58 .69] dva (SS). Other observers' thresholds were .24 [.17 .28] dva (SHK), .34 [.30 .41] dva (CC), .37 [.31 .42] dva (SC), and .59 [.46 .63] dva (KD).

These raw thresholds for each of the transformations specify each observer's sensitivity to that particular shape change. However, it should be noted that these increment thresholds cannot be used to directly compare observers' sensitivity between these transformations because they are in different units (e.g., DVA for length, and degrees for orientation).

A common way to compare sensitivities across different dimensions (even different modalities) is to use Weber fractions, $\Delta I/I$, where I is the baseline intensity and ΔI is the difference threshold. However, computation of Weber fractions is only possible for dimensions defined on a well-defined ratio scale (thus with a unique "0"), which is not the case for orientation. Thus, in order to directly compare the thresholds across all transformations, the thresholds were converted into a common measure of shape difference in the next step of data analysis. (Weber fractions will also be reported for the three transformations for which they can be meaningfully defined.)

2.2.1. Area-based difference measure

For each transformation, an area-based shape difference metric was computed. Given any two shapes, Sh_1 and Sh_2 that are aligned maximally with each other, the metric is defined by the following formula (and depicted schematically in Figure 25):

$$\frac{Area(Sh_1 - Sh_2) + (Sh_2 - Sh_1)}{Area(Sh_1\_part) + Area(Sh_2\_part)}$$

where $Sh_1 - Sh_2$ refers to points of Sh_1 that are not in Sh_2. Thus, Area(Sh_1-Sh_2)+Area(Sh_2- Sh_1) corresponds to the area of the non-overlapping portions of Sh_1 and

Sh_2. The sum of Sh_1 and Sh_2's non-overlapping regions (numerator) is normalized by the sum of the areas of the parts of the two shapes (Area (Sh_1_part) + Area (Sh_2_part)); denominator (Note that the formula normalizes by the sum of part areas, and not by whole shape area in order to facilitate comparison of common metric values between this Experiment and the following Experiment 2, which will test a part-only shape. Moreover given that the transformations only involved the attached part, this ratio cannot exceed 1.) The raw thresholds were converted to this measure by taking the baseline ("test") shape as Sh_1 and the "threshold shape" (i.e., the shape with a particular observer's threshold value added to the baseline value for a given transformation) as Sh_2. The formula produces a proportion (of the part area; which is converted to percentage) which corresponds to the area difference by which the threshold shape differs from the test (baseline) shape.

2.2.1.1. Differential sensitivity across transformations

Figure 26 shows bar plots of each observer's thresholds converted to the common area-based shape difference. From these plots, it is evident that observers are most sensitive to the transformations of part width and length, then to axis curvature, and are least sensitive to transformations involving part orientation and location of attachment.

The average common area metric value across observers for length was 5.12%, for width: 4.23%, for curvature: 14.21%, for orientation: 22.77%, and for location: 44.76%.

Length vs. width. Five of the six observers were more sensitive to the width than to the length transformation; this difference was statistically significant for four of the five observers. One observer was more sensitive to length than to width; this difference was not statistically significant.

Length & width vs. curvature. All six observers were more sensitive to both length and width transformations compared to the curvature transformation. This difference was statistically significant for all six observers.

Curvature vs. orientation. All six observers were more sensitive to curvature than to orientation transformation. This difference was statistically significant for four of the six observers.

Curvature vs. orientation vs. location. All six observers showed a pattern of greatest sensitivity to curvature, lesser sensitivity to orientation, and the least sensitivity to the location transformation. This pattern was statistically significant for three of the six observers.

Orientation vs. location. All six of the observers were more sensitive to orientation, and less to location transformation. This pattern was statistically significant for four of the six observers.

Location vs. all other transformations. All six observers showed the lowest sensitivity to the location transformation compared to the other transformations. This pattern was statistically significant for four of the six observers.

2.2.2. Distance-based difference measure

The goal of converting raw increment thresholds that were obtained from various transformations into a common metric (such as the area-based metric), is to enable meaningful comparisons between sensitivities to different shape transformations despite the fact that thresholds are in different units: in other words, to investigate whether observers exhibit greater sensitivities to some transformations over others.

A natural question is whether the ordering of the sensitivities of different transformations might be somehow due to the specific measure that was used to provide a common scale for all transformations. In order to obtain a greater degree of confidence in this ordering, I used a second, very different, measure based on distance rather than area. This measure is closely related to the Hausdroff measure used in Mathematics and Computer Science. For each point on a given shape (Sh_1), the distance to the closest point on the second shape (Sh_2) is determined. This is done for all points on Sh1, and an average value is calculated, denoted d(Sh1→Sh2). This process is now repeated by beginning with Sh_2: for each point on Sh_2, we find the distance to its closest neighbor on Sh_1, we calculate these distances for all points on Sh_2, and then compute the average value across all points of Sh_2. It should be noted that in general, d(Sh1→Sh2) will not be equal to d(Sh2→Sh1). The average of these two values (Sh_1 to Sh_2, and Sh_2 to Sh_1) is the final distance metric value for any two shapes. The distance-based metric is thus defined by the following formula and illustrated in Figure 27:

$$\frac{d(Sh_1 \rightarrow Sh_2) + d(Sh_2 \rightarrow Sh1)}{2}$$

2.2.2.1. Differential sensitivity across transformations

Figure 28 shows bar plots of each observer's thresholds converted to the distance-based shape metric. The ordering of thresholds for different transformations is very similar to the one obtained with the area-based measure (recall Figure 26): observers are more sensitive to transformations of width, length, width, and curvature than to orientation and location.

The average common distance metric value across observers for length was 0.0362 dva, for width: 0.0303 dva, for curvature: 0.0984 dva, for orientation: 0.1594 dva, and for location: 0.2943 dva.

Length vs. width. Five of the six observers were more sensitive to the width than to the length transformation; this difference was statistically significant for two of the five observers. One observer was more sensitive to length than to width; this difference was not statistically significant.

Length & width vs. curvature. All six observers were more sensitive to both length and width transformations compared to the curvature transformation. This difference was statistically significant for all six observers.

Curvature vs. orientation. All six observers were more sensitive to curvature than to orientation transformation. This difference was statistically significant for four of the six observers.

Curvature vs. orientation vs. location. All six of the observers showed a pattern of greatest sensitivity to curvature, lesser sensitivity to orientation, and the least sensitivity to the location transformation. This pattern was statistically significant for three of the six observers.

Orientation vs. location. All six of the observers were more sensitive to orientation than to location transformation. This pattern was statistically significant for four of the six observers.

Location vs. all other transformations. All six observers showed the lowest sensitivity to the location transformation compared to the other transformations. This pattern was statistically significant for four of the six observers.

2.2.3. Correlation between area- and distance-based metrics across transformations

Figure 29a shows the correlation between area- and distance based metrics, for all increment values (there was a total of seven increment values per each transformation). Visual inspection of the scatter plot shows that an increase in area-based metric is associated with an increase in distance-based metric, indicating a consistent linear relationship between the two metrics. The correlation (Pearson's *r*) was significant for all transformations: length (r=.99, p<.001), width (r=.99, p<.001), curvature (r=.99, p<.001), orientation (r=.99, p<.001), and location (r=.96, p<.01).

Figure 29b shows the correlation between area- and distance-based metrics for the raw thresholds for all six observers; a consistent linearly increasing relationship is observed. Note that the lateral shift seems to be the only transformation that does not fall along the same linear trend as the other transformations; the highest metric values do not bring about increasingly larger differences (compared to lower metric values) because the area-based measure maxes out once the magnitude of lateral shift exceeds the width of the part. The correlation was significant for all transformations: length (r=.99, p<.001), width (r=.99, p<.001), curvature (r=1, p<.001), orientation (r=1, p<.001), and location (r=.98, p<.01). This pattern of the relationships between the area- and distance-based metrics provides confirmation that the differences in sensitivity to shape transformations gauged by two distinct common metrics are reliable. Despite the fact that the two metrics are based on very different measures of shape difference (area-based and distance-based), and in general can yield very different outcomes, in the context of the transformations, they do seem to be highly correlated.

2.3. Discussion

In the two-part shape, observers are most sensitive to the width transformation of the attached part, followed by length, curvature, orientation, and location transformations. This ordering is identical for both the area-based and distance-based metrics as well as for the corresponding Weber fractions (Table 1). Thus, overall, observers are most sensitive to transformations of the shape's part that involve width, length, and curvature, and are least sensitive to those of orientation and location. The shift in location of the part resulted in the poorest sensitivity compared to all other transformations.

What might be the reason for this ordering of sensitivities to different transformations? The transformations of length, width, and curvature involve a judgment about a change involving a single axis of a two-part shape. In contrast, in the *orientation* transformation, observers are asked to make a judgment about the orientation of the part in relation to the base. Note that observers are less sensitive to the orientation change compared to length, width, or curvature.

Sensitivity to *location* change is the poorest among other transformations, indicating that it is harder to notice that the part shifted its location (its point of attachment) along the base. Observers are asked to make a judgment relative to the part's base, and each new location transformation involves a change where the part meets the base. Although the shape's branching topology is the same regardless of the magnitude of shape change, only when the position of the part's attachment to the base changes a great deal do observers notice the change.

The next experiment (Exp. 2) investigates sensitivity to shape transformations of length, width, curvature, and orientation, but now applied to a single-part shape.

3. Experiment 2: Single-axis shapes

The goal of Experiment 2 was to investigate observers' sensitivity to transformations of a simple, single-axis shape. The shape used in Experiment 2 was a stand-alone version of the part from the two-part shape used in Experiment 1. Unlike the composite shapes with two axial branches that were used in Experiment 1, these shapes have a single axis. This makes it possible to directly compare whether, and if so how, the ordering of sensitivities to shape transformations is affected when a given transformation applies to an independent, stand-alone shape, as compared to a part of a composite shape. For example, would observers be most sensitive to transformations that involve a change in length and width of the stand-alone shape, followed by a change in curvature, and least sensitive to a change of orientation (the outcome in Experiment 1 in the context of transformations of one axial branch of the two-part shape)?

Thus the main difference from Experiment 1 is that each transformation was now applied to the single-axis shape. Four types of shape transformations were tested on a simple, single-axis shape: change in the length, width, orientation, and curvature of the shape's axis. The transformation of location was not tested. As in Experiment 1, observers saw the test shape followed by two alternative shapes; all three were masked. One of the alternatives matched the test shape, while the other was modified along one of the transformation dimensions. Using a 2IFC task, observers were asked to make a perceptual judgment: which alternative shape matched the test shape—first or second?

3.1. Method

3.1.1. Observers

The same six observers from Experiment 1 also participated in this experiment.

3.1.2. Apparatus

The equipment was identical to that used in Experiment 1.

3.1.3. Stimuli

An elongated, single-axis shape was used in this experiment. When the axis was straight, the shape was a rectangle with smoothed corners. The dimensions of this shape were exactly the same as those for the attached part of the two-part shape in Experiment 1. Its width subtended .86 degrees of visual angle (dva) and its length subtended 2.58 dva (shape in the 1^{st} column, 1^{st} row: Figure 30). The shape was located near the center of the display. More precisely, it was positioned slightly above (.60 dva) the horizontal meridian and could appear randomly within the range of +/- 1.55 d.v.a. away from the vertical meridian of the monitor on any given trial. This position was identical to the one occupied by the part of the two-part shape (Experiment 1). In addition, in order to avoid judgments strictly at vertical meridian, the rectangle was oriented either + 20 or – 20^{o} away from the vertical meridian.

As in Experiment 1, on every trial observers saw three shapes on the computer monitor: a "test" shape followed by two alternative shapes. The test shape version had identical parameters to the shape described above for transformations of length, width, and orientation. As in Experiment 1, in the curvature condition, the test shape was different in two respects 1: first, it had a baseline orientation of 0^{o} (the orientation of its axis measured at its lower end was vertical; this shape is illustrated in the 1^{st} column, 4^{th} row: Figure 30), and its baseline curvature parameter was set to 0.41 (1/dva) (the axis of the shape was an arc of a circle with the radius equal to 1/ 0.41, i.e., 2.44 dva and could curve either to the right or to the left of the vertical meridian of the monitor). One of the alternative shapes was a transformed version of the test shape while the other was identical to the test shape. The transformations tested were length of the shape, its width, orientation, and curvature.

In the length condition, the alternative shape was different from the test shape in the extent of its elongation. One of seven increments was applied to the initial length parameter of 2.58 dva: .09, .17, .27, .34, .43, .52, & .60 dva. The length transformation is illustrated in the first row of Figure 30.

In the width condition, the length of the ribs (i.e., width) of the shape was modified; one of seven increments was applied to the initial width parameter of .86 dva: .03, .05, .09, .12, .15, .19, & .22 dva. The width transformation is illustrated in the second row of Figure 30.

The orientation condition modified the part's initial tilt away from the baseline orientation of +20 or - 20^{o} from the vertical. One of seven increments was applied to this initial orientation: 2, 6, 10, 14, 18, 22, & 26 degrees. The orientation transformation is illustrated in the third row of Figure 30.

The curvature condition changed the extent of the bending of the shape's axis. One of seven increments was applied to the base curvature parameter of 0.41 1/dva: .06, .12, .17, .23, .29, .35, .41 1/dva (for subjects KD & CC) and .04, .08, .12, .16, .2, .24, .28 1/dva (for subjects SC, SS, SHK), and .01, .05, .09, .13, .17, .22, .26 (for subject SK). The curvature transformation is illustrated in the fourth row of Figure 30.

3.1.4. Procedure

The procedure was identical to that used in Experiment 1. As before, observers viewed three shapes on the monitor in succession: a test shape and two alternative shapes. The task was to indicate which alternative shape matched the test shape, and to press "1" on the keyboard for first alternative and "2" for second alternative (2IFC task). The order of the alternatives (i.e., a transformed version and a version identical to the test shape) was counterbalanced. The observers' responses were not timed, and no feedback was provided after the response. Immediately after the response the fixation cross reappeared to signal the onset of the next trial. The trial sequence was identical to that in Experiment 1 (as illustrated in Figure 23).

3.1.5. Design

The design was the same as in Experiment 1. Each observer participated in a total of 8 experimental sessions, 2 sessions each for the 4 transformations of length, width, orientation, and curvature. Each transformation condition contained a total of 350 trials divided into two sessions (one session contained 140 and the other 210 trials). This resulted in a total of 50 repetitions for each increment value. The order of the conditions was counterbalanced across observers. There was also a brief (20 trials) practice session given to each observer given prior to the experimental session.

3.2. Results

The data was analyzed in the same manner as in Experiment 1. First, Weibull psychometric curves were fit to individual observers' data for each transformation condition using the psignifit software for MATLAB (Wichmann & Hill, 2001a). Based on these fits, increment thresholds and their corresponding 95% confidence intervals were computed. The raw data (proportion correct for each increment level) and the psychometric fits are shown in Figure 31a-d, for each condition. Table 2 shows all the thresholds.

<u>Increment thresholds.</u> Thresholds for length ranged between .14 [.11 .16] dva (observer SK) and .39 [.37 .45] dva (KD). Other observers' thresholds were .16 [.13 .18] dva (SHK), .19 [.15 .21] dva (SC), .21 [.17 .23] dva (SS), and .30 [.27 .34] dva (CC).

Thresholds for width ranged between .07 [.06 .08] dva (observer SS) and .11 [.10 .12] dva (KD). Other observers' thresholds were .07 [.05 .08] dva (SHK), .08 [.07 .09] dva (SC), .09 [.07 .01] dva (CC), and .10 [.09 .10] dva (SK).

Thresholds for curvature ranged between .06 [.05 .08] 1/dva (observer SK) and .15 [.12 .16] 1/dva (KD). Other observers' thresholds were: .06 [.05 .07] 1/dva (SHK), .07 [.06 .08] 1/dva (SS), .08 [.06 .09] 1/dva (SC), and .08 [.06 .10] 1/dva (CC).

Thresholds for orientation ranged between 5.15 [3.99 5.76] degrees (observer SC) and 12.79 [9.99 13.73] degrees (KD). Other observers' thresholds were: 5.76 [4.61 6.8] degrees (SK), 5.92 [4.28 6.87] degrees (SS), 6.15 [4.64 7.17] degrees (CC), and 7.10 [5.76 8.50] degrees (SHK).

3.2.1. Area-based difference measure

As in Experiment 1, thresholds were converted into a common area-based measure of shape difference for the purpose of directly comparing sensitivities across the different part transformations of length, width, curvature, and orientation (illustrated in Figure 32 in the context of a single-axis shape). Given two shapes, Sh_1 and Sh_2, they were first aligned with each other in a way that replicates the alignment of the part in the two-part shapes in Experiment 1. (Additional analysis of both area-based and distance-based metrics where shapes were maximally aligned with each other is presented in a subsequent section) The metric was normalized by the sum of the areas of the two shapes (stand-alone single-axis shapes):

$$\frac{Area(Sh_1 - Sh_2) + (Sh_2 - Sh_1)}{Area(Sh_1) + Area(Sh_2)}$$

3.2.1.1 Differential sensitivity across transformations

Figure 33 shows bar plots of each observer's thresholds converted to the common area-based shape difference. The plots show the same trend in sensitivity to the different part transformations as Experiment 1: observers are most sensitive to changes in length and width, followed by curvature; and are least sensitive to changes in orientation. This trend was statistically significant for five of the six observers. For observer SK, there was no statistically reliable difference between width and curvature.

The average threshold in terms of the area difference measure across observers for length was 4.77%, for width: 5.17%, for curvature: 9%, and for orientation: 14.37%.

Length vs. width. Five of the 6 observers were more sensitive to the length than the width; this difference was significant for 2 observers.

Length & width vs. curvature. All six observers were more sensitive to both length and width transformations compared to the curvature transformation. This difference was statistically significant for two of the six observers.

Length and width vs. orientation. All observers were more sensitive to both length and width compared to orientation; this difference was statistically significant for all six of the observers.

Curvature vs. orientation. All observers were more sensitive to curvature than to orientation; this difference was statistically significant for four of the six observers.

3.2.2. Distance-based difference measure

As in Experiment 1, increment thresholds across different transformations were also compared using the same distance-based measure. Figure 34 shows bar plots of each observer's thresholds converted to the distance-based shape metric.

3.2.2.1 Differential sensitivity across transformations

The pattern of sensitivities is similar to the one obtained with area-based measure: observers are most sensitive to changes in length and width, followed by curvature. Observers are least sensitive to changes in orientation. This trend was statistically significant for five of the observers; for observer SC, there was no statistical difference between orientation and curvature.

The average threshold in terms of the distance-based metric value across observers for length was 0.0299 dva, for width: 0.034 dva, for curvature: 0.1284 dva, and for orientation: 0.0717 dva.

Length vs. width. Four of the 6 observers were more sensitive to the length than the width; this difference was significant for two of the observers.

Length & width vs. curvature. All six observers were more sensitive to both length and width transformations compared to the curvature transformation. This difference was statistically significant for all six observers.

Length and width vs. orientation. All observers were statistically more sensitive to both length and width as compared to orientation.

Curvature vs. orientation. All observers were more sensitive to curvature than orientation. This difference was statistically significant for five of the 6 observers.

3.2.3. Correlation between area- and distance-based metrics across transformations

Figure 35 shows the correlation between raw thresholds converted into area- and distance-based metrics, for all six observers. As for Experiment 1, examination of the scatter plot shows a consistent, linearly increasing relationship between area- and distance-based metrics, indicating agreement between the two distinct metrics. These correlations were significant for all transformations: length (r=.99, $p<.001$), width (r=.99, $p<.001$), curvature (r=1, $p<.001$), and orientation (r=.99, $p<.001$).

3.2.4. Area-based difference measure: shapes are maximally aligned

Recall that the area-based and distance-based metrics for Experiment 1 involved two-part shapes that were first maximally aligned with one another. In order to keep conditions as close as possible to Experiment 1, in Experiment 2 the single-axis shape that was previously attached to base in Experiment 1 was not aligned maximally prior to carrying out the common metric computations; instead, the shapes were aligned just as if they were attached to a larger base. In the current analysis the two shapes, Sh_1 and Sh_2, were maximally aligned according to the shapes' centers of mass (An example is shown in Figure 36). Exactly the same formula and procedure was used as in the previous analyses: thresholds were converted into a common area-based measure of shape difference for the purpose of directly comparing sensitivities across the different part transformations of length, width, curvature, and orientation.

3.2.4.1. Differential sensitivity across transformations

Figure 37 shows bar plots of each observer's thresholds converted to the common area-based shape difference. The plots show a similar trend in sensitivity to the different part transformations as Experiment 2, the non-aligned version, although the differences between transformations are overall less pronounced. Four of the observers are most sensitive to changes in length and width, followed by curvature; and are least sensitive to changes in orientation. The rest of the observers are most sensitive to changes in length and width, followed by orientation, and are least sensitive to changes in curvature. The main overall trend emerged between length and width, as compared to curvature and orientation: observers are most sensitive to changes in length and width compared to curvature and orientation. This trend was statistically significant for four of the six observers. For observer SK there was no statistically reliable difference between width, curvature, and orientation; for observer CC there was no statistically reliable difference between length, curvature, and orientation.

The average threshold in terms of the area difference measure across observers for length was 4.35%, for width: 5.01%, for curvature: 8.43%, and for orientation: 8.53%.

Length vs. width. Four of the 6 observers were more sensitive to the length than the width; this difference was significant for 2 observers.

Length & width vs. curvature. All six observers were more sensitive to both length and width transformations compared to the curvature transformation; this difference was statistically significant for four of the six observers.

Length and width vs. orientation. All observers were more sensitive to both length and width compared to orientation. This difference was statistically significant for five of the six observers.

Curvature vs. orientation. Two of the six observers were more sensitive to curvature than to orientation; however, this difference was not statistically significant.

3.2.5. Distance-based difference measure: shapes are maximally aligned

Increment thresholds across different transformations were also compared using the distance-based measure applied to maximally aligned shapes. Figure 38 shows bar plots of each observer's thresholds converted to the distance-based shape metric.

3.2.5.1. Differential sensitivity across transformations

The pattern of sensitivities is similar to the one obtained with area-based measure: Four of the observers most sensitive to changes in length and width, followed by curvature and then orientation. The rest of the observers are most sensitive the length and width, then to orientation and least sensitive to curvature. Overall, observers are most sensitive to changes in length and width, followed by curvature and orientation. This trend was statistically significant for three of the observers. For observer SK, the difference between width compared to curvature and orientation was not statistically significant; for observer CC, no statistically significant differences were found between the transformations.

The average threshold in terms of the distance-based metric value across observers for length was 0.0299 dva, for width: 0.0342 dva, for curvature: 0.0555 dva, and for orientation: 0.0587 dva.

Length vs. width. Four of the 6 observers were more sensitive to the length than the width; this difference was significant for two of the observers.

Length & width vs. curvature. All six observers were more sensitive to both length and width transformations compared to the curvature transformation. This difference was statistically significant for four of the six observers.

Length and width vs. orientation. All observers were statistically more sensitive to both length and width as compared to orientation. This difference was statistically significant for four of the six observers.

Curvature vs. orientation. Four observers were more sensitive to curvature than orientation. This difference was not statistically significant.

3.2.6. Correlation between area- and distance-based metrics (max. aligned shapes)

Figure 39 shows the correlation between raw thresholds converted into area- and distance-based metrics, for all six observers, for single-axis shapes that were first maximally aligned with one another. As in previous comparisons, examination of the scatter plot shows a consistent, linearly increasing relationship between area- and distance-based metrics. The correlations were significant for all transformations: length

(r=.99, $p<.001$), width (r=.99, $p<.001$), curvature (r=.99, $p<.001$), and orientation (r=.99, $p<.001$).

3.3. Discussion

The results of Experiment 2 showed that observers are least sensitive to changes in orientation of a single-axis shape, when compared with those of length, width, or curvature. The overall pattern of ordering of sensitivities to different transformations is identical to Experiment 1, which involved a two-part shape. It is important to note that the transformations of length, width, and curvature constitute changes to intrinsic properties of the shape in this experiment, whereas the change in shape's orientation pertains to how the shape is embedded in the surrounding space. This may explain why observers are much less sensitive to orientation transformation because in this experiment: it is not intrinsic property of the shape. Note that observers were less sensitive to the orientation change compared to length and width in Experiment 1 as well. (In that experiment, the change in orientation did involve intrinsic shape change.)

Within the intrinsic shape parameters, observers are more sensitive to length and width transformations, than they are to curvature of the axis. Importantly, the ordering of sensitivities (from highest to lowest) is identical not only for both shape-difference metrics (the area-based metric and distance-based metric) but also for corresponding Weber's fractions (presented in Table 2). In summary, this experiment shows that observers are differentially sensitive to transformations of axis length, rib length (related to shape width), and axis curvature of a single-axis shape (no part structure), and these differences are corroborated by three different measures that enable comparisons of sensitivities across different transformations.

3.3.1. Comparison of Experiment 1 and Experiment 2 results

The ordering of common metrics for the transformations of shapes with and without the base is stable across Experiments 1 and 2. However, compared to the part-only shape, observers were less sensitive to transformations of length, orientation, and curvature in Experiment 1, and more sensitive to transformation of width. In Experiment 1, five of the six observers were less sensitive to length transformation of the two-part shape (both metrics). Five of the six observers were also less sensitive to the orientation transformation (both metrics) and five were less sensitive to the curvature change according to area-based metric (all six were less sensitive according to distance-based metric). However, for the width transformation, the opposite pattern is observed: observers were more sensitive to changes involving the one-part shape, compared to a part of the two-part shape. All six observers were more sensitive to width in the two-part shape according to the area-based metric (five of the six were found to be more sensitive if using the distance metric). Thus, although the patterns of sensitivity to different transformations are similar in both experiments, the threshold values were elevated for transformations of length, orientation, and curvature transformations, and diminished for the width transformation, in the two-part shape (Experiment 1) compared to part-only shape (Experiment 2).

The presence of the base part in the two-part shape did not lead observers, overall, to detect shape changes with greater efficiency compared to the single-axis shape. This seems surprising given that, in Experiment 1, the base can serve as a frame of reference for the attached parts, which therefore should make judgments easier than in Experiment 2, where a single-axis shape is involved (especially, for example, in case of orientation transformation). However, this did not happen, since although the overall ordering of sensitivities was similar between the two experiments, thresholds were on average higher, not lower, in Experiment 1.

In summary, a similar pattern of results is obtained when the part that was previously attached to the base (Exp.1) is now shown as an independent, single-axis shape: sensitivity to part elongation and width is best, followed by bending, followed by part articulation (orientation change). An additional condition that was tested in Experiment 1, the shift in the location where part attaches to the base resulted in the worst sensitivity. This type of change is certainly less common in biological (and even manufactured shapes), and is in that sense less natural.

4. Experiment 3: Control for base-part orientation

In Experiment 1 sensitivity to shape transformations was found to be no better than in Experiment 2. Examination of the raw thresholds for most transformations between the two experiments reveals that they are comparable. The average raw threshold for length was .28 dva (Exp1) and .23 dva (Exp2), for width it was .07 dva (Exp.1) and .09 dva (Exp2), for orientation it was 8.93 deg (Exp1) and 7.14 deg (Exp2) and for curvature it was .11 1/dva (Exp1) and .08 1/dva (Exp2). This seems surprising, since the base part of the two-part shape used in Experiment 1 could serve as a frame of reference. In the case of the orientation transformation, for example, observers would only need to make a judgment that involves estimating the angle between the base and the attached part, rather than the absolute orientation of the part on the screen. However, no improvement in sensitivity was observed Experiment 1. One factor that may be relevant is that in Experiment 1 the orientation of the base part was always horizontal. It is possible that, because of its fixed orientation, the presence of the base part may not have been as beneficial. The purpose of Experiment 3 was to investigate whether observers' judgments were affected by the fact that the orientation of the base part in Experiment 1 was fixed.

In this experiment, the orientation of the entire shape (either one-part, or two-part shape) was varied randomly on each trial. The orientation transformation was applied to the single-axis shape from Experiment 2 and two-part shape from Experiment 1. Exactly the same 2IFC task used in the two previous experiments was used (a test shape and two alternative shapes) in which observers were asked to make a perceptual judgment: which alternative shape matched the test shape—first or second?

4.1. Method

4.1.1. Observers

The same six observers from Experiments 1 and 2 also participated in this experiment.

4.1.2. Apparatus

The equipment was identical to that used in the previous two experiments.

4.1.3. Stimuli

Two shapes were used in this experiment: a two-part shape from Experiment 1, and a single-axis shape from Experiment 2. The parameters used in presenting the shapes were the same as in the orientation transformation in the two previous experiments, with the exception that the shapes were presented on the screen at a random orientation on every trial. Specifically, the single-axis shape had no pre-defined baseline orientation (unlike + 20 or – 20^{o} from the vertical meridian in Experiment 2). Similarly, the orientation of the base part in the two-part shape was randomized from trial to trial. However, as in Experiment 1, the part of the two-part shape had a baseline orientation of + 20 or – 20^{o} away from the direction orthogonal to the base of the two-part shape. Similar to the previous experiments, three shapes were presented on each trial: a test shape, followed by two alternatives. In the context of the single-axis shape (“part only”), the orientation transformation involved changing the tilt of the entire alternative shape relative to the orientation of the test shape. For the two-part shape, the tilt of the attached part of the alternative shape was changed relative to the baseline part orientation. Seven increments were applied to the baseline orientation parameter (i.e., +20 or - 20^{o} from the direction orthogonal to the base part orientation in case of the two-part shape): 2, 6, 10, 14, 18, 22, & 26 degrees. The orientation transformation is illustrated in the third row of Figure 30 for the single-axis shape, and in the third row of Figure 22 for two-part shape.

4.1.4. Procedure

The procedure was identical to that used the previous two experiments. As before, observers viewed three shapes on the monitor in succession: a test shape and two alternative shapes. The task was to indicate which alternative shape matched the test shape, and to press “1” on the keyboard for first alternative and “2” for second alternative (2 IFC task). The order of the alternatives (i.e., a transformed version and a version identical to the test shape) was counterbalanced within each experimental session. The observers’ responses were not timed, and no feedback was provided after the response. Immediately after the response the fixation cross reappeared to signal the onset of the next trial. The trial sequence was identical to that in Experiments 1 and 2 and illustrated in Figure 23.

4.1.5. Design

Each observer participated in 4 experimental sessions: two sessions for each of the two conditions (orientation change of a single-axis shape, and orientation change of the attached part of two-part shapes) for a total of 4 sessions. Each transformation condition contained a total of 350 trials divided into two sessions (one session contained

140 trials and the other 210 trials). This resulted in a total of 50 repetitions for each of the seven increment values. The order of the conditions was counterbalanced across observers. There was also a brief practice session (20 trials) given to each observer given prior to the experimental session.

4.2. Results

The data was analyzed in the same manner as the previous experiments. The raw data (proportion correct for each increment level) and the psychometric fits are shown in Figures 40a and 40b, for each condition. Table 3 provides an overview of the thresholds reported in degrees.

Increment thresholds. Thresholds for orientation in the part-only condition ranged between 6.23 [4.63 7.19] (observer SHK) and 21.78 [20.60 24.61] (KD). Other observers' thresholds were 7.32 [5.85 8.86] (SC), 7.64 [6.25 8.77] (SS), 8.17 [6.59 9.42] (SK), and 11.78 [10.22 13.57] (CC).

Thresholds for orientation in the two-part condition ranged between 7.87 [6.32 9.07] (observer SHK) and 13.47 [11.82 15.57] (CC). Other observers' thresholds were 8.22 [6.46 9.36] (SK), 10.80 [9.22 12.48] (SC), 11.99 [10.59 13.70] (SS), and 12.82 [11.34 14.93] (KD).

4.2.1. Differential sensitivity between single-axis and two-part shapes.

Since this experiment involves a single transformation type, namely change in orientation, it is meaningful to directly compare the raw increment thresholds between the two conditions (single-axis shape and two-part shape). The average raw threshold (across all 6 observers) for the single-axis shape was 10.48 degrees, and for the two-part shape, 10.86 degrees. Three of the six observers' raw thresholds were significantly different between the two conditions; of these, two observers' thresholds were higher for the two-part shape, and for one observer the threshold was lower for the two-part shape. The comparison between the two types of shapes presented at random orientations indicates that there is no evidence of improvement in sensitivity to orientation change in the case of the two-part shapes.

4.3. Discussion

Sensitivity to shape transformations was no better in Experiment 1 than in Experiment 2. This was not expected because the base in the two-part shape provides a frame of reference for observers' judgments that is not available in the single-axis shape, and therefore, should have yielded better sensitivities for shape transformations. Was this lack of benefit simply because the orientation of the base part in Experiment 1 was always fixed (i.e., always positioned horizontally with respect to the display)? Experiment 3 indicates that even when the orientation of the entire shape was changed randomly from trial to trial, the sensitivity still is no better in the two-part shape than in the single-axis shape. In other words, consistent with the comparison between Experiment 2 and Experiment 1, the presence of the base part does not appear to improve sensitivity (even though it can potentially provide a strong frame of reference for the orientation of the attached part).

5. Experiment 4: Positive vs. negative part transformations

Experiments 1, 2, and 3 examined observers' sensitivity to transformations involving either a single-axis shape or a two-part shape that had a part attached to its base, and found that the sensitivity (as measured by shape-difference metrics that provide a common unit for all transformations) varies according to the type of transformation for each of the two shapes. The goal of Experiment 4 was to investigate how surface geometry (as opposed to contour geometry) influences sensitivity to shape transformations.

The influence of contour- vs. surface-based representation has often been investigated empirically by manipulating figure and ground. By keeping the contour between two regions fixed, but making one or the other side figural, researchers can alter the geometry of the perceived surface. When this occurs, for example, convexities become concavities, and vice versa. There is a great deal of evidence indicating that this switch affects shape perception. For example, observers were able to judge the position of vertices that were perceived as convex faster than those that were perceived as concave (e.g., Bertamini, 2001). A series of studies utilizing a visual search paradigm found superiority for finding targets containing concavities within convex distractors compared to finding targets with convexities among concave distractors (e.g., Hullman et al., 2000; Humphreys and Muller, 2000). In the context of change detection, observers were found to be more sensitive to a change involving a concave vertex. Specifically, sensitivity was higher with an addition or a modification of a concave, compared to a convex vertex of a polygon (e.g., Barenholtz et al., 2003; Cohen et al., 2005). Furthermore, in the context of figure and ground assignment, Barenholtz & Feldman (2006) found an "articulating concavity bias" such that observers assign figure when a portion of the display articulates at concavity.

Recent empirical studies also investigated sensitivity to concave compared to convex figures. For example, Bertamini & Farrant (2005) found higher sensitivity for concave vertices when part structure of the shape is affected. On the other hand, Bertamini et al. (2008) found that observers' responses were faster when a convex surface was presented as figure ("convex-in-front") while it controlled for the amount of surface area between the convex and concave surfaces. When two surface regions have equivalent areas, observers are foveating the convex (positive) part surface but not the complementary convex (negative) part surface. A follow-up experiment by Bertamini et al. (2008) manipulated the extent of the convex and concave surfaces but nonetheless found the same advantage for detecting a surface bound by convex contours. Thus, the issue of which kinds of surfaces the visual system is more sensitive to (and which types of changes are allocated more representation in the visual system) has not been completely resolved.

In the current experiment, the contour geometry of the display was kept constant, while the surface geometry was manipulated by changing figure-ground relations using binocular disparity. The figure-ground relationship is demonstrated schematically in Figure 41. As illustrated, the two regions, black and white, share the same central undulating contour. When the black region is designated as figure, it appears as a shape with a part protruding out from its base (i.e., a positive part); the white region is seen as a background. When figure and ground assignment is reversed, and the white region is designated as figure, it appears as a shape that has a cavity (i.e., a negative part) within its base. This shape has a very different surface geometry and axial structure compared to the shape with the protrusion, even though the shared bounding contour remains the same.

In addition, unlike the previous experiments in which the entire bounding contour of the experimental shape was displayed, only the critical portion of the shape was visible through a circular aperture (the entirety of the shape was not shown). The rationale for this approach was motivated by a pilot study, which showed differential sensitivity to positive and negative part transformations. The positive and negative parts were manipulated by enclosure (i.e., the shape was enclosed on all sides; see Figure 42). The shape with a positive part was similar to that in Experiment 1 (with the small exception that its part was attached to a wider base part), and the other was its complement (a shape with a deep indentation in its part) (Figure 42, top row). Results of this pilot study indicated that observers are more sensitive (i.e., they require less shape change) when a transformation (e.g., of length of the part) involves a negative rather than the positive part. However, one concern with that study was that subjects may be making a judgment based on the negative part's proximity to the nearest edge of the base part. For example, changing the length of the negative part will cause the indentation to deepen and approach the opposite contour of the shape's base (Figure 42, bottom row). If this is the case, then this portion of the contour can be used as a frame of reference, and any measured threshold reflects not only sensitivity to the transformation manipulated, but also proximity to such landmarks (subjects could be focusing on the gap between the tip of the indentation and bottom portion of the shape's contour.) The current experiment addresses this concern by showing shapes through a circular window which is placed closer in depth than the shape (neither side of the shape is enclosed); observers never see the entire bounding contour of the shape (unlike Experiments 1, 2, and 3).

To summarize, the *bounding contour* of the part of the shape used in this experiment is identical to that used in Experiment 2, with the exception that the bulge of the contour could face either to the left or to the right. The regions on the either side of the contour constituted of complementary figure and ground in the display. This figure and ground relationship reversed systematically in accordance with the experimental design. Specifically, the contour remained the same while figure and ground relationship was reversed; this allowed a single contour, (e.g., with bulge facing to the right) to produce the percept of either a positive or a negative part (Figure 43). Instead of aligning the longer axis of the shape along the horizontal meridian of the monitor as in Experiment 2, the orientation of the base contour was aligned along one of two diagonals: either +20 or -20^{o} away from the horizontal meridian. In addition, a total of two types of bulging contours were presented, one with its "bulge" facing to the left, and one to the right. (The exact configuration of different types of stimuli are described in detail in a separate section below).

Two part transformations, orientation and location change, were tested; both were presented as positive and negative versions. Note that these are the two transformations that involve the spatial relationships between both part and the base. The orientation transformation represents a shape change that normally (i.e., naturally) occurs in objects, especially part articulation in biological objects. An example of such shape change is moving one's arm at the shoulder. The second transformation, the shift in the part's location relative the base, is unnatural and observers do not seem to be very good at detecting this type of change.

Each part transformation was presented as one of two types: either as a positive-part change or as a negative-part change. One objective of the current experiment was to examine whether observers are more sensitive to either positive or negative-part changes. For example, the prediction that observers may be more sensitive to negative part changes would be consistent with the finding that changes occurring at concavities are easier to detect (e.g., Barenholtz et al., 2003; Cohen et al., 2005). A second objective is to examine whether there is a pattern of interaction of the type of part (i.e., positive or negative) with the type of transformation (i.e., orientation or location change). For example, it would be interesting to examine whether the direction of the pattern of sensitivity between positive and negative transformations for orientation condition is the same as for the location condition.

5.1. Method

5.1.1. Observers

Six observers participated in this experiment: three observers who had participated in Experiments 1 and 2, plus three were new observers. All were experienced psychophysical observers. All observers had normal or corrected-to-normal vision, and screened for stereoscopic vision.

5.1.2. Apparatus

Stimuli were generated using MATLAB (Mathworks) and the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997). Stimuli were presented on a high resolution (1280x1024) 19-inch monitor (Mitsubishi DiamondPro) connected to a dual-core G5 Macintosh computer with a 100 Hz refresh rate. The stereo images were alternately presented to the right and left eyes using an infrared emitter and stereoscopic shutter glasses system StereoGraphics CrystalEyes (described in detail below).

<u>StereoGraphics CrystalEyes</u>. In order to induce stereoscopic viewing, this system uses goggles with liquid crystal shutters that are synchronized with the stimulus presentation on the CRT monitor. An infrared (IR) emitter placed on top of the monitor sends IR signal pulses that are sensed by the goggles, thus controlling the occlusion or transmission of the light that enters the goggles, separately for either the left or the right eye. Briefly, the shutters in the goggles are comprised of three layers: the liquid crystal cell and a sheet of polarized material on either side of the cell. Since each sheet is coated with a kind of transparent conductor (indium tin oxide), an electric field can be set up between the two sheets upon application of voltage. As Lipton (2001) describes, "when a field is induced, the liquid crystals become isotropic; thus "crossed" sheet polarizers block light from the display monitor". When no field is applied, the light is allowed to enter the goggles and observer perceives displayed images.

This manipulation permits either the right or the left eye to perceive the stimulus display meant for that eye. The visual system integrates the two images into a unified display; the images are constructed such that they differ slightly in horizontal displacement (i.e., binocular displarity) in order to induce depth perception. The alternation of the images sent to the left and the right eye reduced in half the effective refresh rate for each eye. Specifically, this resulted in a refresh rate of 50 Hz per each eye. The coordination between Stereographics emitter graphics, and the computer was implemented using a ‘stereo mode option’ in Psychtoolbox.

5.1.3. Stimuli

Stereo presentation. Random Dot Stereograms (Julesz, 1971) were used to present stimulus displays. Each “dot” was a square of 4x4 pixels (4.2 x 4.2 arc min) with equal probability of black and white dots. The entire stereo display was presented within a square measuring 900 by 900 pixels (15.47 x 15.47 dva) against a black background. The trial sequence is shown schematically in Figure 44. Every trial began with a frame displaying the fixation cross. The following frame displayed a circular window through which the background surface was visible, but no shape was presented. The next several frames contained the three experimental shapes that were shown through the circular window: a test shape followed by two alternative shapes; all three were masked. The mask frame consisted of the same circular window shown in all of the frames presented with a different background (i.e., one of 16 different randomly generated dot patterns as background was presented). Thus the frames in the trial sequences (except the first frame which was the fixation cross) were of two types: (1) those that consisted of the circular window and the background without an experimental shape and (2) those that displayed the circular window and the background along with one of the experimental shapes.

Depth relationships: Frames without experimental shapes. The depth relationship between the circular window and the background is illustrated in Figure 45a: the circular window appears closer in depth relative to the background. For the circular window, a new random dot pattern was generated at random for each trial. For the background (seen through the window), a random dot pattern was selected randomly for every display frame of the experimental trial from one of 16 different, previously randomly generated dot patterns as background; this effectively serves as a mask when this type of display follows a frame that contains a shape. The disparity between the circular window and the background surface was 15.6 min arc.

Depth relationships: Frames with experimental shapes. The depth ordering of the circular window, the experimental shape, and the background surface is illustrated in Figures 45b and c: the circular window appears closest in depth, the background appears farthest, and the experimental shape is positioned between the window and the background. The random dot pattern for the circular window was the same as that on the first frame of the trial; the dot pattern for the background was generated in the same manner as for frames when no experimental shape is present (i.e., randomly chosen from the 16 pre-generated patterns), and the dot pattern for the shape itself was generated at random for every frame within the trial. The disparity between the circular window and the shape was 5.4 min arc, while the disparity between the shape and the background surface was 10.32 min arc.

Shapes. A total of four shapes were used. The shapes were defined by one of two contours: one whose bulge faced to the left of the vertical meridian of the display, and one whose bulge faced to the right of the vertical meridian. Since the same contour gives rise to two differently shaped surfaces, depending on figure and ground relationships, two different shapes resulted from each of the two contours (see Figure 43). Specifically, when given a right-facing bulge, and when the region to the right of the contour was presented as figure, it produced a shape with indentation (Shape #2, Figure 43), but when the region to the left of the contour was presented as figure it produced a shape with a protrusion (Shape #1, Figure 43). Given a left-facing bulge, when the region to the left of the contour was presented as figure, it produced a shape with an indentation on its right side (Shape #4, Figure 43). On the other hand, when the region to the right of the contour was presented as figure, it produced a shape with a protrusion on its left side (Shape #3, Figure 43). Thus, there were a total of 4 possible combinations (figure type (positive or negative) x bulge direction (bulge facing left or facing right)). In summary, the visible portion of all shapes included the part that either protruded from the middle of the shape's base, or the part that constituted an indentation into the shape. The rest of the shape was occluded behind the circular window; no information regarding the extent of the shape's continuation behind the window was available.

Shape Transformations. The transformations of orientation and location were applied to the part of the shape. The orientation transformation involved increasing the magnitude of the part's initial tilt of +20 or - 20^{o}. The positive-part transformation modified the tilt of the protrusion (positive part), and negative-part transformation modified the tilt of the indentation (negative part). One of seven increments was applied to this initial orientation: 2, 6, 10, 14, 18, 22, & 26 degrees. Orientation transformation is illustrated in the third row of Figure 22.

Location transformation shifted the location of the part relative to the base. The positive part transformation shifted the location of the protrusion (positive part), and negative part transformation shifted the location of the indentation (negative part). One of seven increments was applied to the initial location of the part relative to the horizontal center of the shape (either +1.20 or –1.20 dva): .14, .28, .42, .55, .69, .83, or .96 dva.

Location transformation is illustrated in the fifth row of Figure 22. The entire stereo display subtended 15.47 by 15.47 dva.

Stereo test. Since the task in this experiment relied on observers' ability to discriminate depth, before running in experimental trials all observers were screened for stereo vision using a custom-developed stereo test. Observers were sequentially presented with one of 6 shapes that were not part of the experiment: a vertical rectangle, a horizontal rectangle, a vertical ellipse, a horizontal ellipse, an upward pointing isosceles triangle, and a side-ways pointing isosceles triangle. The shapes were presented against a uniform gray background, and given one of two disparities: crossed, or uncrossed. The magnitude of disparity was 8.4 arc min for both types of disparities. In the crossed disparity condition, the shape appeared floating in front of the background surface, while in the uncrossed disparity, it appeared as a hole with that shape through which another textured surface was visible. Each observer was shown 12 trials (6 shapes x 2 disparities). Each shape was presented for 1 second each. The observers were required to make three responses: (1) Triangle, Rectangle, or Oval? (2) Horizontal or Vertical? (3) Front or Window? For all three questions, observers were instructed to indicate which of the three shapes they saw, and the experimenter entered the response using the computer keyboard. Responses were recorded, and the following criterion was used for judging whether the observer possessed appropriate visual ability to see stereo displays: if observer attained around 92% correct (11 out of 12 trials responded correctly) for all three questions, he or she was allowed to proceed to the stereo experiment. All six observers showed performance at or above 92% correct (most were at 100% accuracy) and thus allowed to proceed to the experimental session.

5.1.4. Procedure

Every trial began with a fixation cross. The first frame shown was a circular window through which a background was visible but no shape was presented. Next, three shapes were shown sequentially through a circular window: a test shape followed by two alternative shapes. All three shapes were masked. (The mask frame consisted of the same circular window shown in all of the frames, but a different random-dot background: one of 16 different randomly generated dot patterns as background was presented). The trial sequence is shown in Figure 44. The task is identical to the previous three experiments: in a 2IFC task, observers were asked to indicate which interval contained the alternative that matched the test shape.

5.1.5. Design

Each observer participated in 10 experimental sessions, 5 sessions for each of the transformations of orientation and location. Each transformation had a total of 700 trials (350 for positive and 350 for negative part transformations). Positive and negative part transformations were interleaved in a single experimental session. One session contained 140 trials and was divided by a break in the middle. This resulted in a total of 50 repetitions for each of the seven increment values for negative and positive transformations. The order of the conditions was counterbalanced across subjects. Each observer completed one practice session prior to the experimental sessions (140 trials).

5.2. Results

The data was analyzed in the same manner as in the previous experiments (Weibull psychometric curves were fit to individual observers' data for each transformation condition and difference thresholds and corresponding 95% confidence intervals were computed). The raw data (proportion correct for each increment level) and the psychometric fits are shown in Figure 46a and b, for each condition. Table 4 provides an overview of the thresholds, reported in dva for the location condition and in degrees for orientation condition.

Increment thresholds. Thresholds for positive orientation condition ranged between 3.63 [2.49 4.33] degrees (observer HH) and 7.2 [5.93 8.17] degrees (MZ). Other observers' thresholds were 4.88 [3.81 5.50] degrees (HH), 5.31 [4.24 6.10] degrees (SHK), 6.11 [4.90 8.53] degrees (CC), and 7.03 [5.64 8.21] degrees (AH).

Thresholds for negative orientation condition ranged between 3.89 [2.66 4.56] degrees (observer SK) and 14.65 [11.29 21.09] degrees (CC). Other observers' thresholds were 4.18 [3.01 5.34] degrees (AH), 5.15 [3.59 5.95] degrees (SHK), 6.93 [5.58 8.24] degrees (HH), and 10.10 [8.86 11.71] degrees (MZ).

Thresholds for positive location condition ranged between .21 [.16 .25] dva (observer HH) and .41 [.37 .47] dva (AH). Other observers' thresholds were .24 [.14 .29] dva (CC), .32 [.27 .36] dva (SHK), .35 [.30 .39] dva (MZ), .35 [.29 .38] dva (SK).

Thresholds for negative location condition ranged between .18 [.03 .27] dva (observers CC) and .41 [.37 .47] dva (AH). Other observers' thresholds were .27 [.22 .33] dva (MZ), .29 [.25 .32] dva (HH), .30 [.25 .33] dva (SK), and .35 [.29 .40] dva (SHK).

Since the main interest in this experiment is comparing negative vs. positive part transformations, conversion to a common metric (area-based and distance-based) was not necessary. Figure 47 shows bar plots of each observer's thresholds for positive and negative part changes for orientation and location conditions.

5.2.1. Differential sensitivity for Orientation: positive vs. negative part transformations

The average threshold for positive part was 5.69 [4.51 6.81] degrees, and for negative part, 7.48 [5.83 9.46] degrees. Thus, on average, observers were more sensitive to the positive-part transformation than to the negative-part transformation. Four of the six observers showed a significant difference between positive and negative transformations. Three observers were more sensitive to the positive transformation, and one was more sensitive to the negative transformation.

5.2.2. Differential sensitivity for Location: positive vs. negative part transformations

The average threshold for location transformation's positive part was .32 [.26 .37] dva, and for the negative part, .30 [.24 .36] dva. Five of the six observers exhibited no significant difference between positive and negative part transformations. The sixth observer was more sensitive to the positive part transformation.

Thus, the results show some evidence for greater sensitivity to positive-part transformation over negative-part transformation in the case of orientation. However, no such difference is evident in the case of location change.

5.3. Discussion

The goal of Experiment 4 was to investigate the extent to which shape geometry (specifically, contour-based and surface-based aspects) has an influence on observers' sensitivity to shape transformations. The shapes' surface geometry was manipulated while the contour geometry was kept the same. The overarching question was whether sensitivity to shape transformations is different for shapes that have identical bounding contours, but different, complementary surface regions (via figure-ground manipulation).

Thus a given contour (such as that in Figure 43) can give rise to two surfaces: one with positive curvature and thus a positive part, and the other with an indentation corresponding to a negative part. In the current experiment, observers' sensitivity to detect changes in orientation and location of a positive or a negative part of a two-part shape was tested.

For orientation transformation, there was some evidence sensitivity was higher to positive than negative part changes, indicating that observers may be less sensitive to changes involving the negative part. What could be a reason for the presence of differential sensitivity between the positive and negative orientation transformations, but not for positive and negative location transformations?

The reason that the thresholds are lower for positive part orientation (compared to the negative change condition) may be due to fact that orientation is a common and natural transformation (i.e., articulating parts) for positive parts, but not for negative parts. Therefore, in the case of the orientation transformation, there is a great difference between sensitivity to positive compared to negative part transformations. In contrast, location change (i.e., lateral shift) is *not* a common or natural transformation—neither for positive nor for negative part. Consequently, there is no reason to expect any difference in sensitivity between positive and negative part location.

The differential sensitivity to positive compared to negative part orientation changes is also consistent with the articulating concavity bias. A positive part transformation involves a modification of the part's axis in relation to the base since the positive part "articulates" at a concavity—at the negative minimum. Thus the heightened sensitivity to positive part changes may be due to modification occurring at, or near, the shape's concave vertex (e.g., Barenholtz et al., 2003; Cohen et al., 2005; Barenholtz & Feldman, 2006).

A greater surface area is allocated for the positive part transformation compared to the negative part transformation. Thus a possible cause for heightened sensitivity could be that observers are more simply more sensitive to a type of change that involves larger overall area, regardless of the bounding contour (in this case, a convex bounding contour is used for positive part transformation). Since previous research does not show a clear pattern regarding the surface vs. boundary debate, even when the surface area is fixed, more investigation of this question is needed (e.g., Bertamini et al., 2008). However, it does not seem that that surface area *per se* is responsible for the difference in sensitivity in Experiment 4: no difference in sensitivity was found for location transformation, indicating that the *kind* of part change affects perceptual sensitivity.

The positive part transformation also involves a change in position of the convex tip of the positive part. Throughout the experimental trial, observers were instructed to keep their eyes at the center of the display (although eye movements could not be fully precluded). At least some portions of the display, especially those containing the tip of the positive part (but not the negative part) fell onto the peripheral portion of the retina containing magnocellular cells sensitive to motion. Since the tip of the convex, positive part is a salient portion of the shape, the trials where the shape change was large produced an abrupt 'motion'. Thus the heightened sensitivity to positive part changes could also be due to the position of the part in relation to the display.

6. General Discussion

6.1. Axial and part-based representations of shape

Previous research indicates that the human visual system represents shapes in terms of simpler components, or parts, and spatial relationships between those parts (Biederman, 1987; Hoffman & Richards, 1984; Marr & Nishihara, 1978). Part-based representations are robust to a variety of configurations that parts of a single, especially biological object, can take. How can the underlying structure of the parts of a shape be represented? One approach is that the parts of a shape can be represented using the underlying, corresponding skeletal structure (Blum, 1973). A shape's skeletal representation can include its main axis and a number of axial branches representing each of the shape's parts. The skeleton allows for a compact, yet flexible representation because it preserves qualitative information about the shape including the location of the parts relative to the base and their branching structure. Importantly, the original shape can be easily reconstructed based on axial-structure information. For example, studies that investigated the effects of changing the number of parts of an object (or axial branches) by adding or removing a concavity found that observers are more sensitive to such changes (e.g., Barenholtz et al., 2003), compared to sensitivity to changes (addition or removal of convexity) that do not affect the qualitative part structure of the shape. However, previous work has not systematically investigated shape transformations that preserve qualitative axial/part structure, while manipulating specific quantitative parameters in an axial representation.

6.2. Approach & experimental strategy in addressing these questions

The overarching question for the current project was, which types of shape transformations is the visual system more sensitive to? The current experiments applied shape transformations involving quantitative variables in a shape's axial representation (e.g., part length, width, orientation, curvature, and location) that do not alter the qualitative axial or part structure, and compared sensitivities to these different transformations. The goal was to examine to whether certain shape transformations, although mathematically (geometrically) equivalent in terms of the magnitude of change, are perceptually different from one another. Recent studies utilized the "change detection" methodology to measure whether the visual system is able to detect certain kinds of shape changes more effectively than other types of changes. In a typical study, observers are shown two shapes, and asked to judge whether they are the same or different from one another (on some trials, there is no change, and on other trials, the two shapes indeed differ). Performance is measured in terms of accuracy of response as a function of different types of changes. For example, in previous studies, shape change was measured in tasks that measured observers' performance to detect either a convex or concave change (e.g., Cohen et al., 2005). This method is helpful for exploring which types of changes are more salient for the visual system, but it does not give a precise measure of exactly how much more change is needed to detect shape change, nor is it possible to compare the precise magnitudes of these amounts between the various types of transformations.

The current study investigated observers' sensitivity to different types of shape transformations by using a 2IFC task and measuring increment thresholds. Thus for each shape transformation, the threshold indicates precisely how much change is needed to perceive that one shape is noticeably different from another shape. Different shape transformations were investigated, such as change of a part's length, width, orientation, curvature or location in a two-part shape and in a single-axis shape (without part structure). Once the difference thresholds were obtained for each of the transformations, they were converted into common units based on shape-difference metrics (area-based and distance-based). The two metrics are tools that allow us to compare the visual system's sensitivity across different types of shape transformations.

6.2.1. Experiment 1

In Experiment 1, several shape transformations were applied to the small attached part of a two-part shape. Transformations of length, width, curvature, orientation, and location were applied to the shape's part, which was attached to a larger base. The two common metrics, area-based and distance-based metrics, were used to gauge sensitivity to different transformations in common units. Area-based metric showed that observers are overall most sensitive to transformations of width and length, followed by curvature, orientation, and location. Similarly, the distance metric showed that observers are most sensitive to width, followed by the length, curvature, orientation, and location. Importantly, the ordering of sensitivities was the same for these two different metrics, indicating that the measured differences between the sensitivities to different transformations are objective and reliable.

The three transformations that involve solely the attached part (a single axis) of the two-part shape—the width, length, and curvature transformations, are more easily detected than those where a judgment involves the relationship between the attached part and the base part of the shape, such as orientation and location. In the orientation transformation, the judgment about orientation requires a comparison across the attached part and the base part. The change in location of the part was the worst (most difficult to detect a change) among all five transformations in the two-part shape. Although the overall branching topology of the shape remains the same, the location transformation does not preserve the location of the part boundaries where the attached part meets the base part of the shape (the location of the part boundaries is preserved in the orientation condition).

6.2.2. Experiment 2

Experiment 2 investigated observers' sensitivity to these same transformations (with exception of location, which was not used) in a simple, single-axis, elongated shape: a change in length, width, curvature, and orientation. Therefore it is possible to directly compare how the ordering of sensitivities to shape transformations may be affected when a given transformation involves an independent, stand-alone shape compared to a part of a complex shape. Examination of the area-based metric values showed that observers were most sensitive to changes in length followed by width, and then curvature; and are least sensitive to changes in orientation. Average distance metric revealed the same outcome: observers were most sensitive to changes in length, followed by width and then curvature and were least sensitive to changes in orientation. Importantly, the ordering of sensitivities was the same for these two different metrics, indicating that the measured differences between the sensitivities to different transformations are objective and reliable. The differences in sensitivity between transformations of length, width, and curvature compared to orientation can be attributed to the fact that the first three transformations involve a change of the intrinsic property of the shape—either its width, length, or the amount of bending of the axis. The transformation of the orientation in this case involves a judgment about the shape's placement in the surrounding context. As in Experiment 1, both metrics' orderings for the different transformations are identical.

While the preceding analysis preserved the shape alignment conditions in Experiment 1, an additional analysis was also conducted in which the shapes were first maximally aligned with one another by their centers of mass for the purpose investigating whether the ordering of sensitivities would be comparable under these conditions. Two common metrics, area-based and distance-based measures were computed using the same procedure as in the previous analysis. Examination of both area-based and distance based metric showed that observers were most sensitive to changes in length and width, and least sensitive to changes in curvature and orientation. Although observers were overall least sensitive to orientation, the difference between curvature and orientation conditions was not statistically significant.

It is important to note that the ordering of common metric values is stable across transformations of a two-part shape (Experiment 1) and a single-axis shape (Experiment 2). For both experiments, sensitivity to length and width was best, followed by curvature and then orientation and location (Experiment 1).

6.2.3. Experiment 3

In Experiment 1, shape transformations were applied to the attached part of the shape. The presence of the base could serve as a frame of reference and facilitate observer judgments such that it would be expected that the thresholds are lower for the two-part shape than those for the single-axis shape. However, thresholds were not found to be lower in Experiment 1. Experiment 3 was conducted to examine whether the fixed overall orientation of experimental shapes in Experiments 1 and 2 affected observers' sensitivity to detect change. One shape transformation, orientation, was tested for both a single-axis and the two-part shape. Unlike in Experiments 1 and 2, however, the overall orientation of the test shape varied randomly on each trial. The raw thresholds for the two kinds of shapes were comparable; in fact, two of the observers were significantly better in the single-shape compared to two-part shape. As in Experiments 1 and 2, sensitivity was no better for the two-part shape compared to the single-axis shape, despite the fact that the base part could potentially provide a frame of reference frame for the orientation of the attached part.

6.2.4. Experiment 4

The goal of Experiment 4 was to investigate the extent to which different aspects of shape geometry (specifically, contour-based and surface-based aspects) influence sensitivity to shape transformations. In this experiment, the contour geometry of the display was kept constant, while the surface geometry was manipulated using binocular disparity in stereoscopically presented figure-ground displays. Under these experimental conditions, it is possible to examine sensitivity to shape transformations that are defined by identical contour but complementary surface areas. Two contours with "bulges" facing left or right produced two shapes with either a positive or a negative part. Two transformations, orientation and location change, were tested.

For orientation transformation, there was an overall difference between positive and negative part changes; however, no difference was observed for location transformation. Higher sensitivity to positive part orientation changes suggests a preference to detect changes that are ecologically valid since the relationship between the positive part (axial branch) and the base of the shape corresponds to part articulation in biological shapes (e.g., such as when an arm swings at an elbow). However, negative changes are less likely to correspond to ordinary shape transformations, and thus may have lower predictive weight for the visual system. On the other hand, location transformation (either a positive or a negative part changes) does not approximate a typical relationship between parts of a complex shape because changing location of the part effectively changes where one axial branch (part) is attached to the larger base (i.e., not normally observed in common objects).

This outcome of this experiment shows that the visual system encodes not only specific parts of an object, but specifically the relationships (e.g., place of attachment) between the part and the shape. Not all "relationships" have equal importance or weight, however, because in the orientation transformations, the location where part attaches to the base is preserved, only the orientation (tilt) is changing. In the location transformations, the location where the part attaches to the base changes.

6.2.5. The role of memory in psychophysical perceptual tasks

It is important to consider one additional issue in discussing the results of these experiments: the role that the working memory may have in change detection tasks whose goal is to measure purely perceptual influence of experimental attributes. Upon first receiving sensory information, for example, through eyesight, the stimulus is said to be in the "sensory register" for the first 600 ms to 1 second (Palmer, 2001, p. 645). Short-term (or working) memory is thought to be involved on the order of several seconds to approximately one minute; it is capable of holding about four items. Visual long-term memory lasts on the order of minutes and days, and even longer (months and years). Most psychophysical experiments reviewed in this thesis last on the order of about one second or more (e.g., Cohen, Barenholtz, Singh, & Feldman, 2005: 900ms; Cohen & Singh, 2006: 1500ms) the current experimental studies are somewhat longer (e.g., Experiment 1-3: 2,700 ms, & Experiment 4: 4,200 ms). Therefore, one question that can arise when interpreting the results of the experimental trials last longer than approximately half a second is, what role does memory play in (mostly) perceptual outcomes in such experiments?

In experimental tasks that require observers to make a comparison between the stimulus shown previously and the current stimulus (e.g., change detection tasks), the first stimulus that was presented to the observer must be retained in memory in order to perform well on the this task. Specifically, there is physiological evidence illustrating that visual features are maintained in memory in neural areas as early as V1 (Supèr et al., 2001). For example, Supèr et al. (2001) observed "contextual modulation of activation in V1" for test stimulus during *delay* periods showing that relevant information is actively stored for later utilization, even when no test stimulus is present in the display.

Previous psychophysical studies also investigated the relationship between perception and memory. For instance, Melcher (2006) found that visual information can accrue over the course of the experimental trials while viewing a stimulus such as a complex natural scene. In this study observers were presented with natural scenes containing various objects for 5, 10, or 20 seconds. Some trials were immediately followed by a memory test (such as reporting color of the item in the scene); those scenes that were not immediately tested were repeated during the experimental session, and a memory test administered. Results showed that performance increased linearly with longer total time views, supporting a theory of memory "accumulation". Observers did significantly better on those trials that were repeated than those that were not, indicating that information about features in the scene is sustained in the working memory even when there are intervening other scenes (or distractors lasting as long as 60 seconds: Experiment 2) before the original scene is repeated.

In addition, the influence of differential task demands (i.e., memory load) during a change detection task has been investigated by Droll & Hayhoe (2007). Observers were asked to perform an object-sorting task while also monitoring for a change in that object. Results showed that when the change involved an object feature that was relevant to the sorting task, observers were significantly better at detecting the change. These results corroborate the physiological findings (e.g., Supèr et al., 2001) that show superior encoding and processing of relevant information in visual working memory. Thus, it is important to note that perceptual sensitivity to shape transformations measured in the current experiments may also reflect some of these modulations at the early stages of visual processing, since the initial test shape needed to be maintained in working memory for the duration of the entire trial.

6.3. Main objective of the study

This series of experiments investigated observers' sensitivity to different types of shape transformations. Changes were applied to two-part shape, a single-axis shape, as well as to shapes whose contours were identical but the surface shape was manipulated (positive vs. negative parts, Exp.4). Observers were sensitive to those transformations that represent common and natural changes in biological objects (orientation transformation, Exp.4); but results were less clear when the change was not natural (location transformation, Exp. 4).

6.4. Qualitative vs. quantitative shape transformations

It is well established that given two shape transformations, one that significantly changes the shape's part structure (i.e., a new part is introduced) is easier to notice compared to one that does alter part. For example, studies that manipulated the spatial configuration of features (i.e., such that an object's features occupy new positions) and substitution of features (i.e., a change in a feature's identity) found superior performance to detect a change in parts' configuration compared to detection of a part's identity (e.g., Keane, Hayward, Burke, 2003; Kimchi & Bloch, 1998; Simons & Levin, 1997). Both theoretical and empirical approaches of shape representation provide support for the notion that such qualitative changes (e.g., changes to the qualitative branching structure) are much more dramatic and thus noticeable than a quantitative change. Current theories of part-based, structural shape representation postulate that the visual system decomposes visual input into smaller units, such as pre-defined shape primitives (geons; Biederman, 1985; 1987), or that the parts of a shape are extracted following a part segmentation process at negative minima (Hoffman & Richards, 1984). The shape description is captured by an underlying axis-based (e.g., medial-axis) representation (e.g., Blum, 1973, Feldman & Singh, 2006; Kovacs & Julez, 1994; Kovacs, Feher, & Julez, 1998; Leyton, 1988; 1989). The critical contribution of these theories is that they capture the relationship between the parts of the shape (parts' configuration), distinguishing such relationship from any particular features or the identity of the parts, and thus transformations of parts of the shape can be used to predict the types of changes expected or not expected by the visual system.

Very little is known empirically about the visual system's sensitivity to transformations within an axial framework: those shape changes that do not qualitatively alter the shape's fundamental relationships between axial branches. The current study investigated perceptual sensitivity to shape changes that occur along a particular axis (simple single-axis shape) or an axial branch (two-part shape) in order to constrain the qualitative geometry (i.e., axial and part structure) while varying quantitative parameters.

This series of experiments demonstrated significant differential sensitivity to shape transformations that change only the quantitative parameters of a single shape without altering that shape's fundamental, qualitative structure. Specifically, the visual system is differentially sensitive to the different types of transformations since each required different magnitude of change to reach the same level of perceptual sensitivity. The differential sensitivity was found in transformations involving solely the intrinsic parameters of a single-axis shape, or an axial branch of the two-part shape. Observers were significantly more sensitive to length and width transformations compared to the curvature transformation.

Those transformations that involved a relationship between two axial branches—such as orientation and location—required a greater amount of shape change to be noticeable. Why is sensitivity different to transformations that required use of two axial branches, as opposed to a single axial branch? Previous studies observed that perceptual judgments are facilitated when a task involves a single object, rather than distinct object (the object superiority effect; e.g., Duncan, 1984; Watson & Kramer, 1999). The variant of the object superiority effect, the *part-superiority* effect, was investigated by Barenholtz & Feldman (2003) who controlled for the magnitude curvature of the local segments involved in the judgments and found that judgments were indeed faster in the "within-parts" condition (i.e., a part was bounded by negative minima) compared to a "between-parts" condition (i.e., when the judgment involved a portion of the shape that was bound by positive maxima).

6.5. Influence of contour- vs. surface- based aspects on shape representation

This study also examined the role of contour-based vs. surface-based properties on sensitivity to different shape transformations. The fact that information carried by the shape's bounding contour plays a role in shape representation is documented by empirical studies, both neuropsychological and psychophysical. As described earlier, neurophysiological evidence indicates that distinct sections of cells in V4 are selectively tuned to curvature: either segments of negative or positive curvature, as well as the magnitude of curvature (Pasupathy & Connor, 1999; 2001; 2002). In another study, observers identified contour segments with negative-minima boundaries with greater accuracy than segments with positive-maxima boundaries (Cohen & Singh, 2007).

Recent studies indicate that shape representation is not solely contour-based, but also region-based (Blum, 1973; Feldman & Singh, 2006; Sebastian & Kimia, 2005). In Experiment 4, when the shape's bounding contour geometry was fixed, but the surface area was changing via a figure vs. ground manipulation, significantly different sensitivities emerged for the two types of transformations. Differential sensitivity between positive and negative part changes in the orientation transformation indicates the role of region-based, compared to solely contour-based, aspects in shape representation since sensitivity was better for those transformations that mimic natural object articulations (positive part changes) compared to those that do not (negative part changes). Sensitivity is better neither for positive nor for negative part in the location condition.

Clearly, both contour- and region-based aspects are important in shape representation, their relative importance may vary depending, in part, on the class of shapes. For example, in the context of shape matching, Sebastian & Kimia (2005) tested shape matching algorithms using curves (contour) or shock-graphs (medial axis-based) and found that shock-graph based matching was more robust (thus more appropriate) for shapes that have part structure, since it captures relationships and hierarchical structure among regions within a shape (e.g., parts), compared to those that involve solely the contour (e.g., handwritten letters). The authors also note that the shock-graph-based algorithms are computationally more intensive.

6.6. How sensitivity to transformations informs similarity measures

Analysis of shape, specifically, the development of robust shape similarity algorithms, is an area of extensive research in computer vision. Such algorithms allow recognizing two objects as being similar to one another although they may not be necessarily geometrically (mathematically) equivalent, and have been utilized in a wide variety of image processing applications, including facial recognition and pattern recognition of highly variable classes of objects.

However, the majority of existing approaches that measure shape similarity are based on a physical measure, not a perceptual one. The results of these current studies provide perceptual information that could potentially be used by such algorithms. The current experiment shows that the human visual system is differentially sensitive to different transformations of shapes that have same qualitative underlying axial structure. For example, the location transformation has very poor sensitivity associated with it: one can change the location parameter a great deal, but the two shapes will still be perceptually similar (i.e., they will have a very strong degree of similarity). On the contrary, a transformation of length or width demonstrates a different degree of sensitivity, since even a very small change is noticeable. This indicates that the visual system differentially sensitive to different dimensions along which shapes can transform: highly sensitive to the length, width, and less so to curvature transformation, and much less sensitive to orientation and location. Such information may thus be relevant to computer vision researchers interested in developing measures of shape similarity that are informed by humans' perceptual sensitivity to shape transformations.

Figures

Figure 1a. Mach (1914/1959) squares illustration. The two squares on the left and right are geometrically identical. When the square on the left is tilted 90 degrees, however, it is perceived as a diamond shape rather than a square (Mach, 1914/1959).

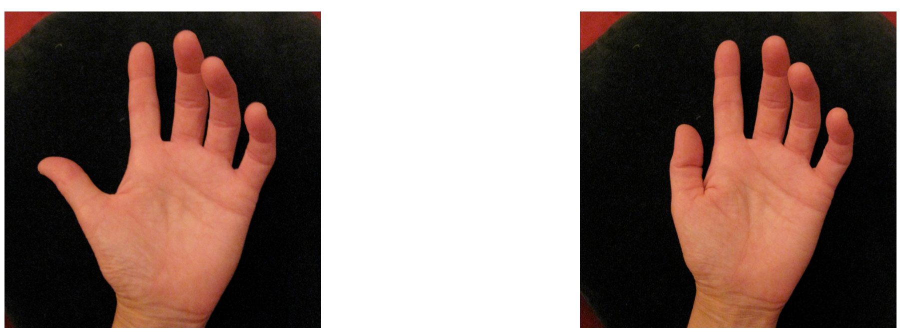

Figure 1b. Articulating hands. The left and the right images depict the *left* hand in two configurations: with the thumb outstretched (left image) and with the thumb close to the other fingers (right image). Both images are geometrically distinct and yet can be perceived to be equivalent.

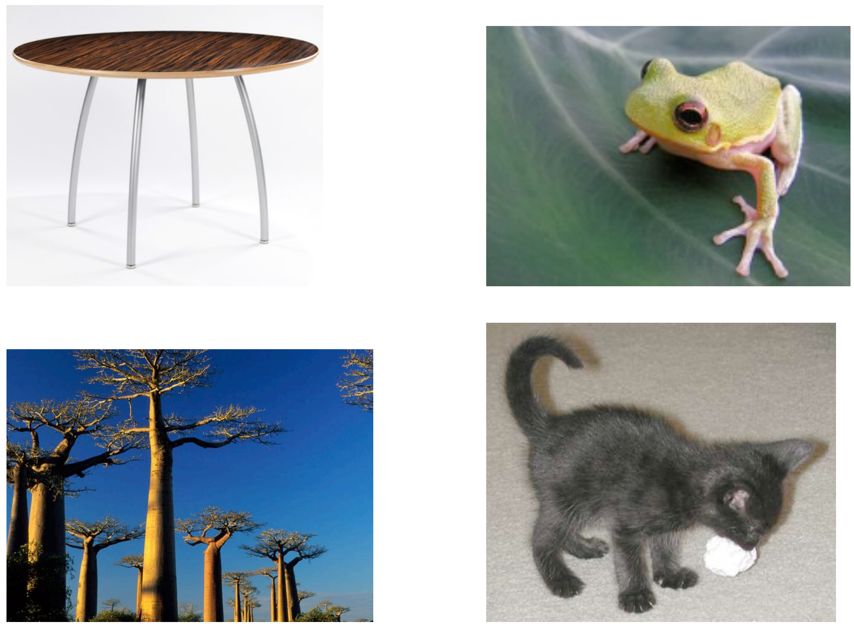

Figure 2. We normally perceive objects as consisting of distinct parts. The table on the left is seen as a round top supported by four legs, the frog is seen as having a trunk and legs (and bulging eyes on the head), the tree is seen with the trunk and branches and the kitten with a curly tail.

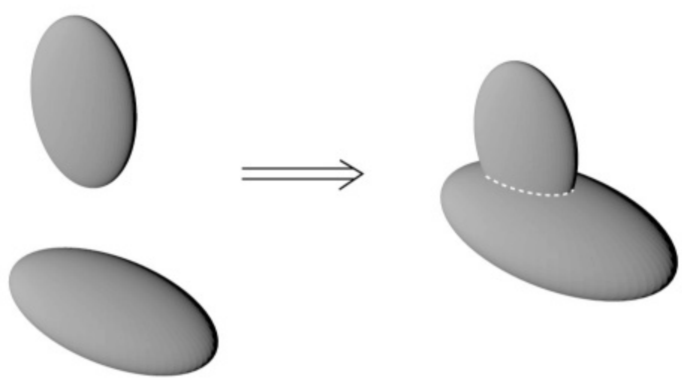

Figure 3. Transversality principle. When any two objects intersect, a concave tangent discontinuity is formed at their intersection (Hoffman & Richards, 1984).

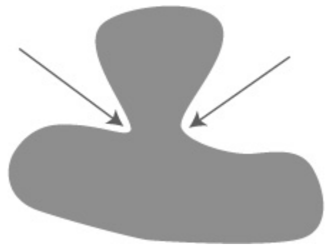

Figure 4. Illustration of negative minima. Smoothing of the concave discontinuities produces negative minima of curvature (Hoffman & Richards, 1984).

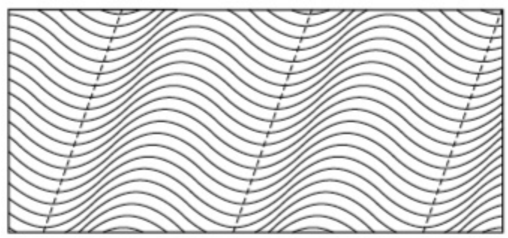

Figure 5. Hills and Valleys. Depending on the viewpoint, observers perceive the dotted line lying inside a valley or, when the figure is turned around 180 degrees, the dotted line lying on top of a hill (Illustration: Singh & Hoffman, 2001).

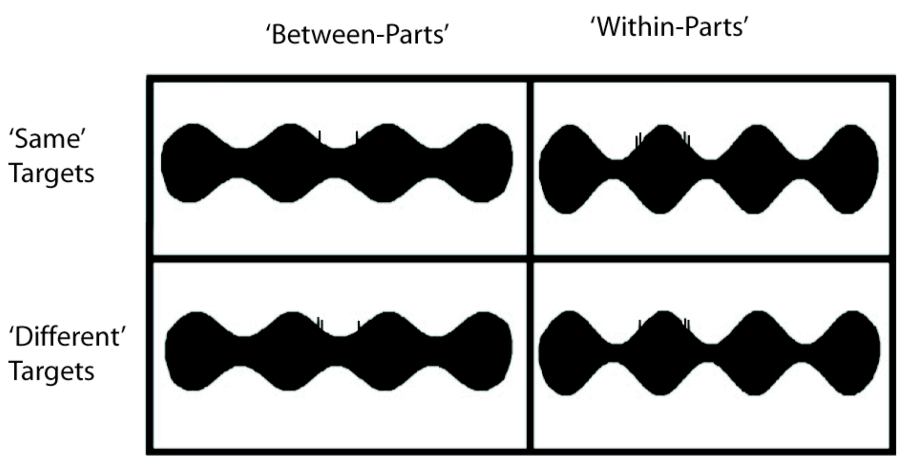

Figure 6. Illustration of experimental stimuli in Barenholtz & Feldman (2003) study. Elongated shapes were used as stimuli which had sets of marks placed along the contour. The marks were placed so that either a negative minimum or a positive maximum extrema was between the two marks. The two marks consisted of either 1 or 2 lines (Barenholtz & Feldman, 2003).

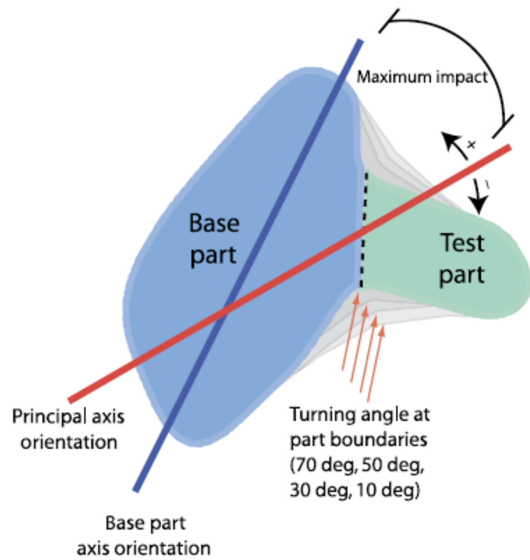

Figure 7. Illustration of the experimental shape (Cohen & Singh, 2006). The experimental shape consisted of a part attached to a larger base. The strength of the part boundary was manipulated using the turning angle near the part's boundaries (Cohen & Singh, 2006).

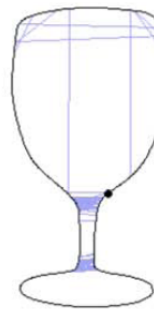

Figure 8. Example of a segmented object, a wineglass. As can be seen in the illustration, most of the segmentation lines fall around the base and top of the stem (DeWinter & Wagemans, 2006).

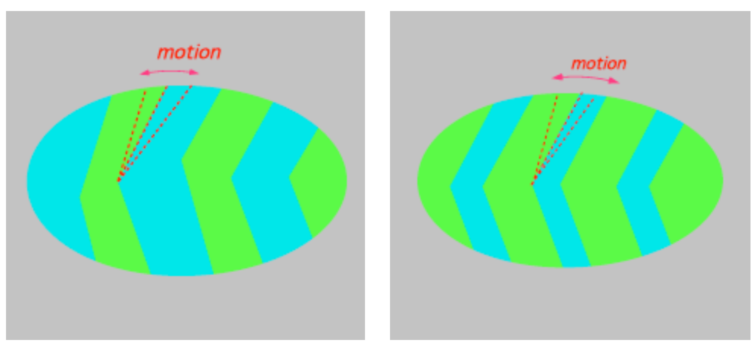

(a) (b)

Figure 9. Articulating concavity cue. When the display in illustration *a* (Experiment 1) and *b* (Experiment 2) is static, the display can be perceived either as green stripes on a blue background, or blue stripes on a green background. However, when the vertex is set in motion (in the illustration, a green stripe moving back-and-forth indicated by the dotted lines) observers judge that articulating concavity as belonging to figure (Barenholtz & Feldman, 2006).

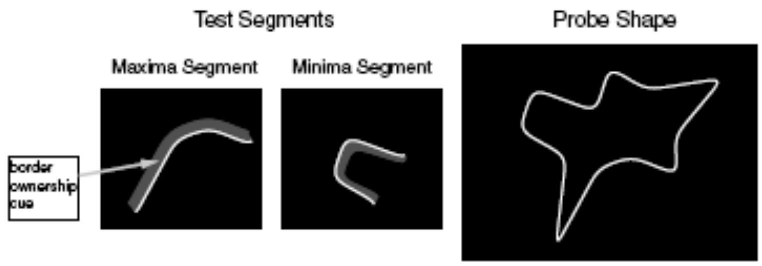

Figure 10. Illustration of experimental stimuli in Cohen & Singh (2007). Observers were shown a probe shape, followed by a test segment. The sample maxima and minima segments are illustrated; the shaded border cue indicated the “inside” of the probe shape (Cohen & Singh, 2007).

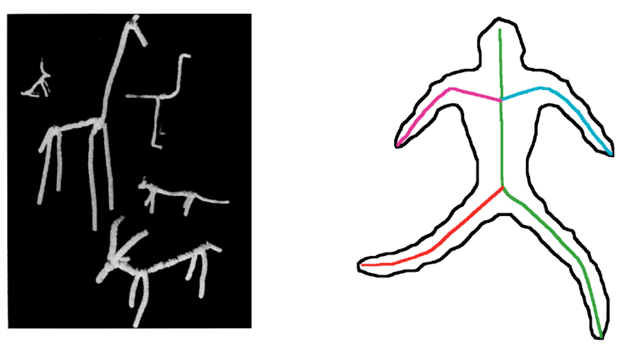

Figure 11. Illustration of a skeletal representation. A “pipecleaner” object is shown in the left image; a skeletal representation is shown in the right image (Marr & Nishihara, 1978; Feldman & Singh, 2006).

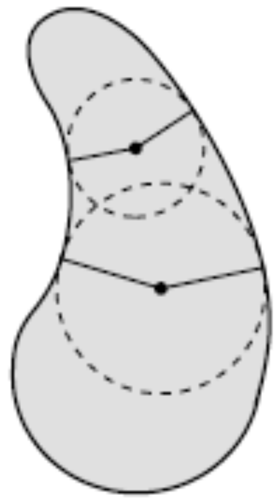

Figure 12. Illustration of Blum's (1973) Symmetry Axis Transform. Medial disks are inscribed within a shape; axis is constructed by connecting the center points of each of the disks (Illustration: from Singh & Hoffman, 2001).

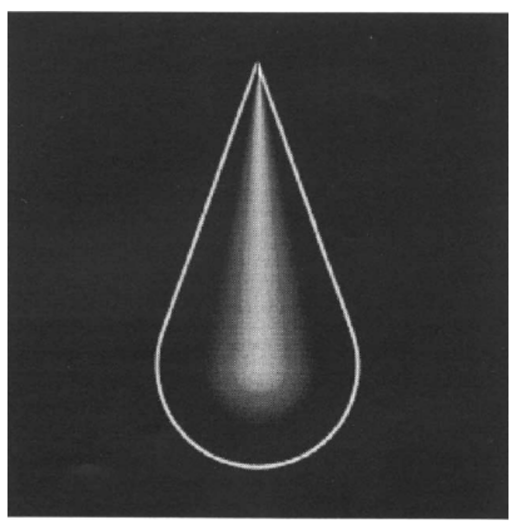

Figure 13. Core analysis. A teardrop figure can be represented by its core. Note that the width of the core is greater towards the "base", corresponding to the increase in width of the actual object (Burbeck and Pizer, 1995).

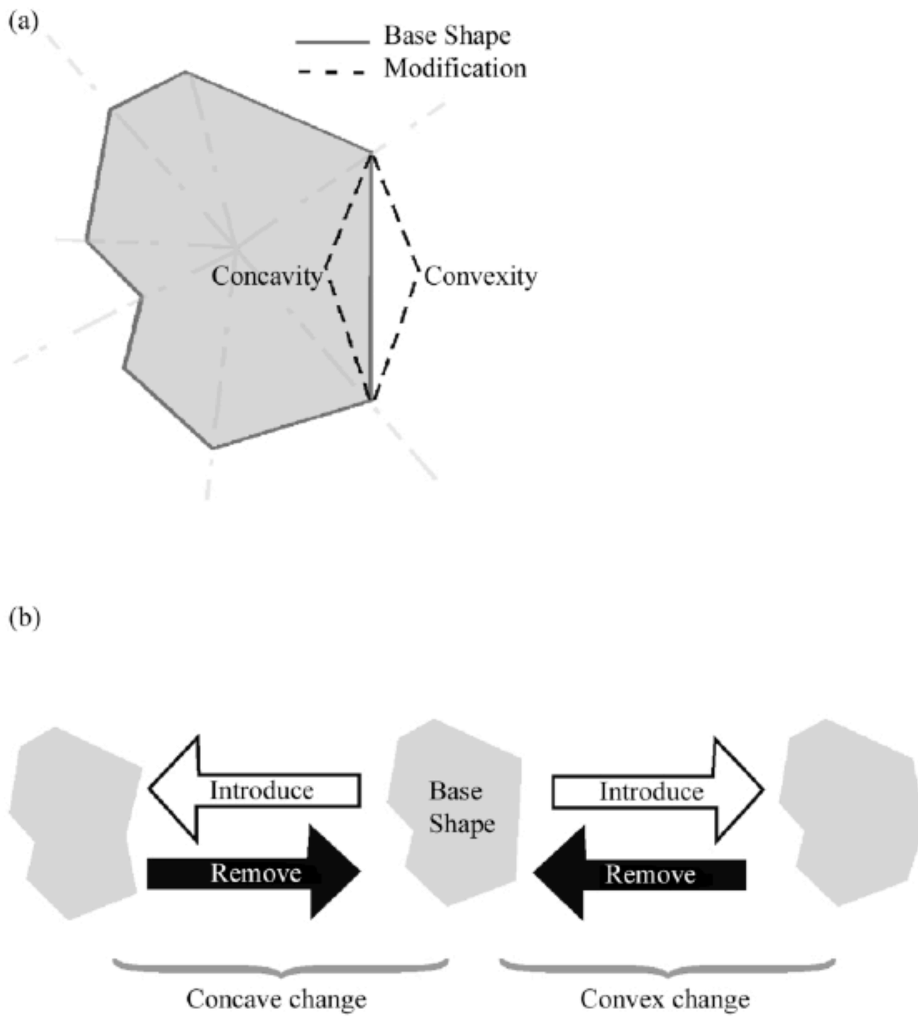

Figure 14. Illustration of experimental stimuli in Barenholtz et al. (2003). Given a base polygon shape (top, (a)), either a concave (left panel, (b)) or a convex (right panel, (b)) could be introduced or removed from that base shape. (Note that the overall shape was thus modified by addition or subtraction of the vertex) (Barenholtz, Cohen, Feldman, & Singh, 2003).

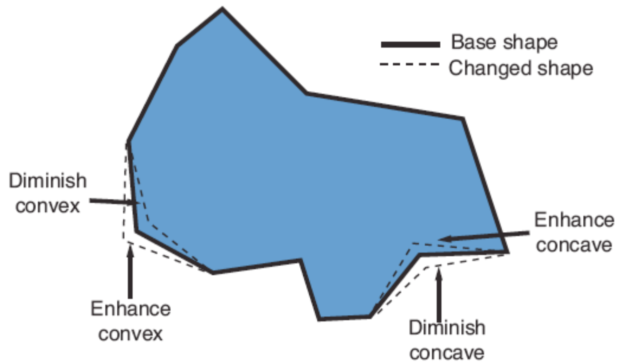

Figure 15. Experimental manipulation in Cohen et al. (2005). Experiment 1: base shape was changed by enhancing or diminishing an existing concavity or convexity (note that the overall shape structure was not modified) (Cohen, Barenholtz, Singh, & Feldman, 2005).

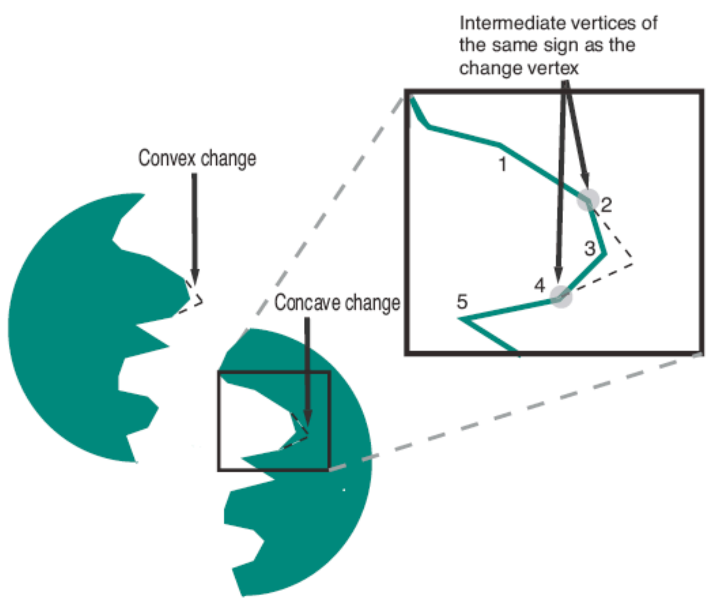

Figure 16. Reversal of figure and ground: Atteneave's egg design (Experiment 2). Two experimental shapes shared the same contour; each shape has complementary part structure, so that a concavity on one half of Atteneave's egg is convexity on the other half (Cohen, Barenholtz, Singh, & Feldman, 2005).

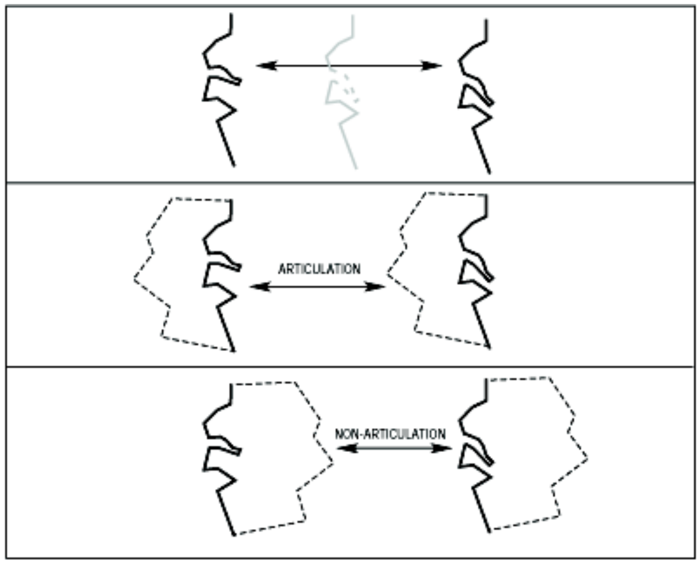

Figure 17. Figure and ground assignment leads to a biologically valid or invalid articulation. The top panel illustrates a portion of a contour rotating at curvature extrema. The middle panel illustrates the case when the figure is completed to the left (with contour to the right), the resulting part is convex and articulation occurs at concavities. The bottom panel illustrates the case when the figure is completed to the right (with contour to the left), and the resulting part is concave (i.e., “negative part”), now rotating at convexities. The middle panel demonstrates a biologically valid articulation, while the bottom panel demonstrates a biologically invalid articulation (Barenholtz & Tarr, 2008).

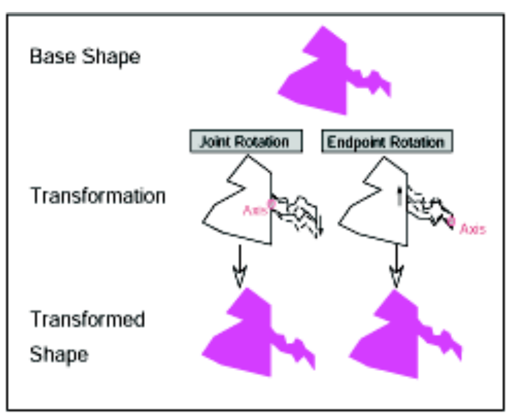

Figure 18. Axis-of-rotation reversal illustration. A given shape’s part can either articulate such that it preserves the spatial relationship of that part to the larger base (left portion of the figure: joint rotation), or such that it does not preserve that spatial relationship (right portion: endpoint rotation) (Barenholtz & Tarr, 2008).

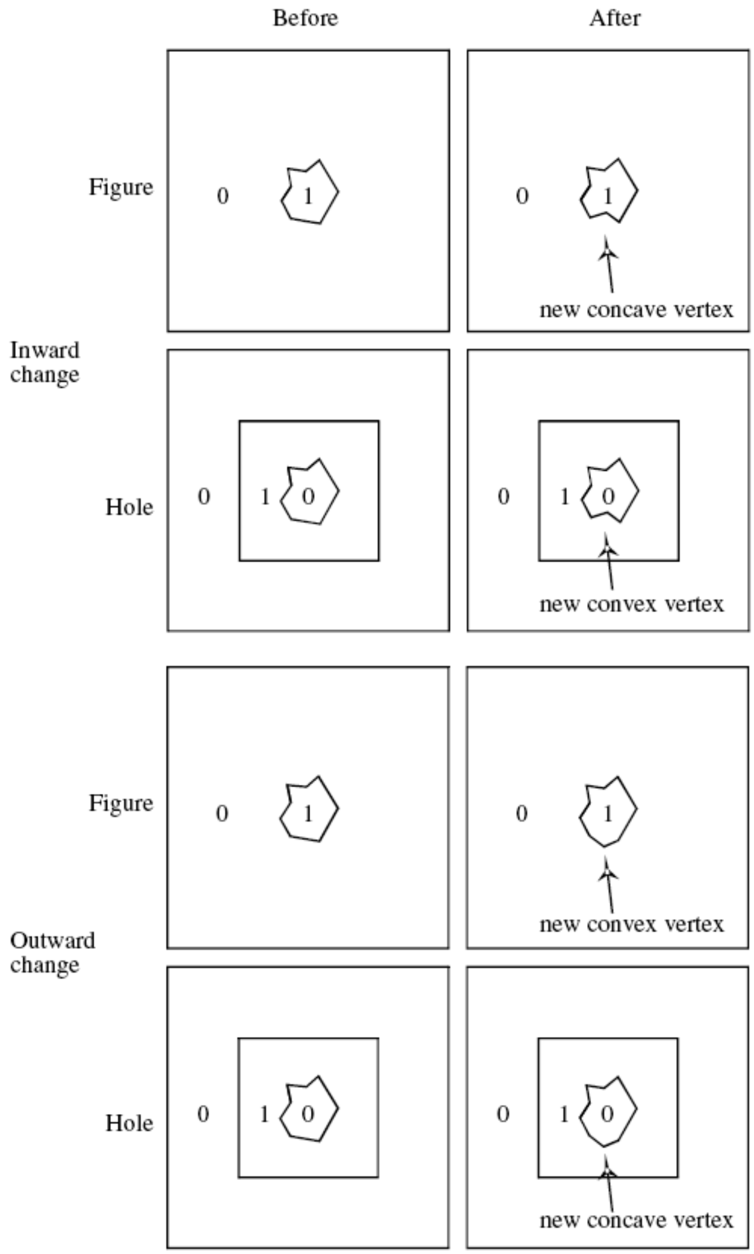

Figure 19. Illustration of experimental stimuli in Bertamini & Farrant (2005). In Experiment 1, direction of vertex change was either inwards or outwards, and was applied either to a figure or a hole. Figure and ground relationship is designated as "1" indicating crossed disparity, and "0" uncrossed (Bertamini & Farrant, 2005). Sensitivity to detect shape changes varied along with the type of vertex change.

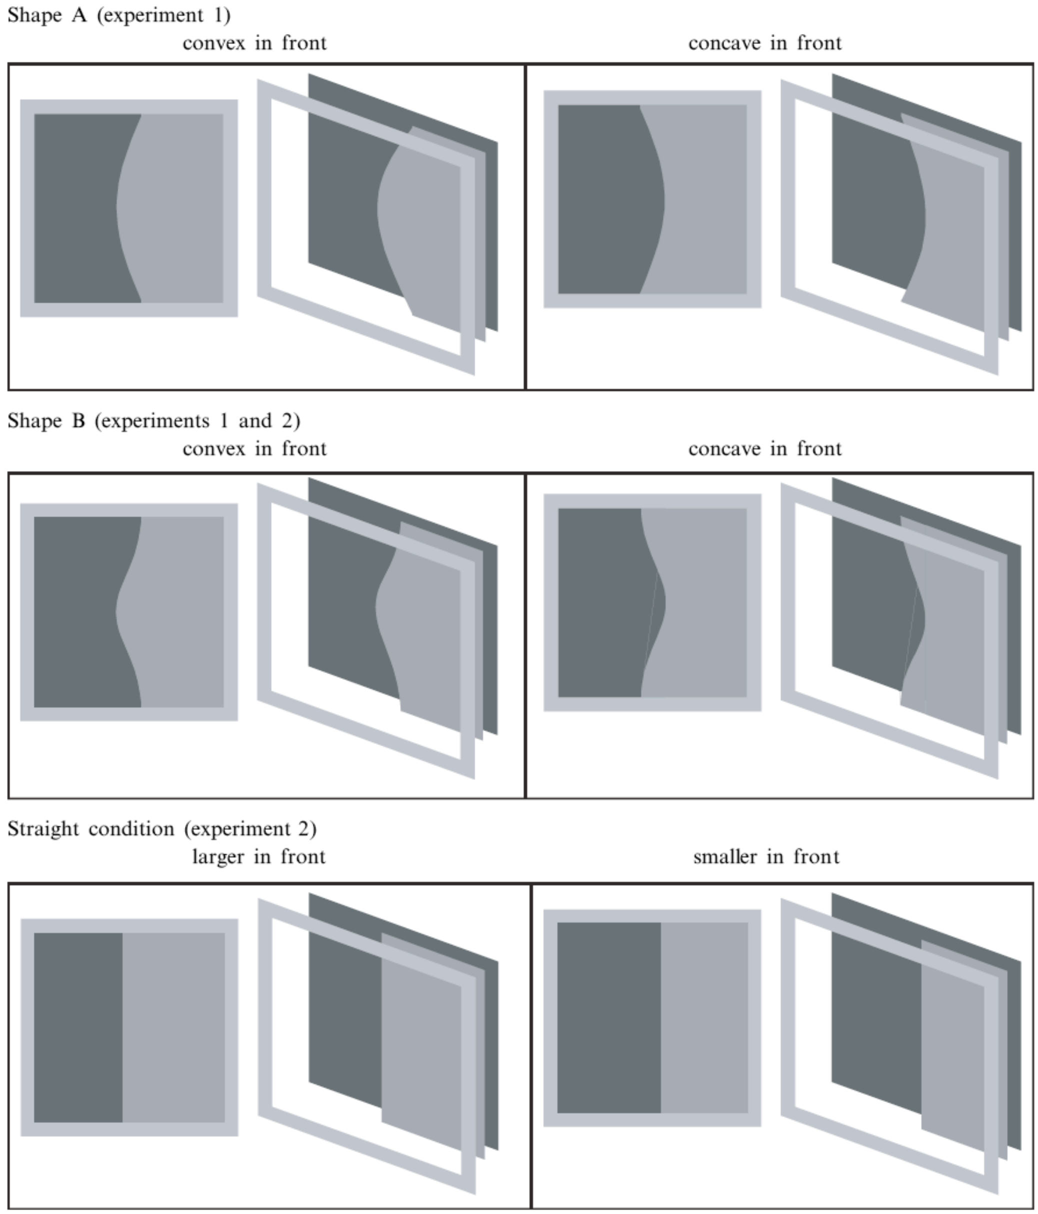

Figure 20. Illustration of shapes A and B used in Experiments 1 and 2 (Bertamini & Lawson, 2008). Top panel illustrates shape A (convex in front: left; concave in front: right). Middle panel illustrates shape B (convex in front: left; concave in front: right). Bottom panel illustrates the straight condition, in which either a larger or smaller area was in front (Bertamini & Lawson, 2008).

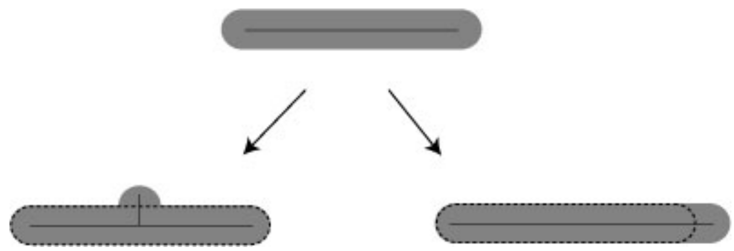

Figure 21. Relationship between sensitivity and similarity among transformations of a simple shape. Transformations that involve a comparable amount of physical change can be perceived as more or less distinct than the original shape. A base shape (an elongated tube) is shown with two transformations below; the base shape's outline is shown in dotted line superimposed on each transformation. Given the same base shape, the transformations on the left and right look distinct. The shape on the left has a small bump that makes it perceptually very different from the base shape. The shape on the right (which has increased in its length) looks similar to the base shape even though, physically speaking, it is more different in terms of area non-overlap, for example.

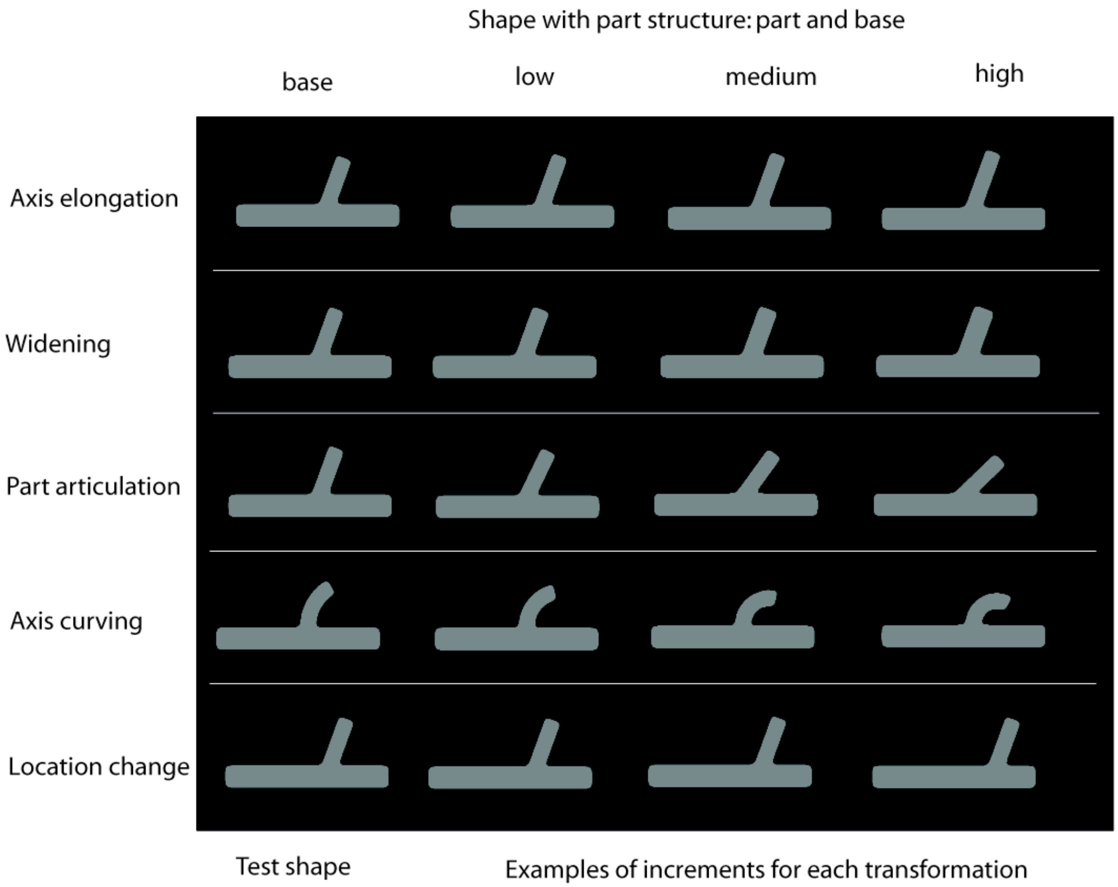

Figure 22. Experimental stimuli used in Experiment 1. Five transformations are shown for the two-part shape: a shape that has a part attached to its base: axis elongation (length), width (widening), orientation (part articulation), curvature (axis curving), and location change of the part.

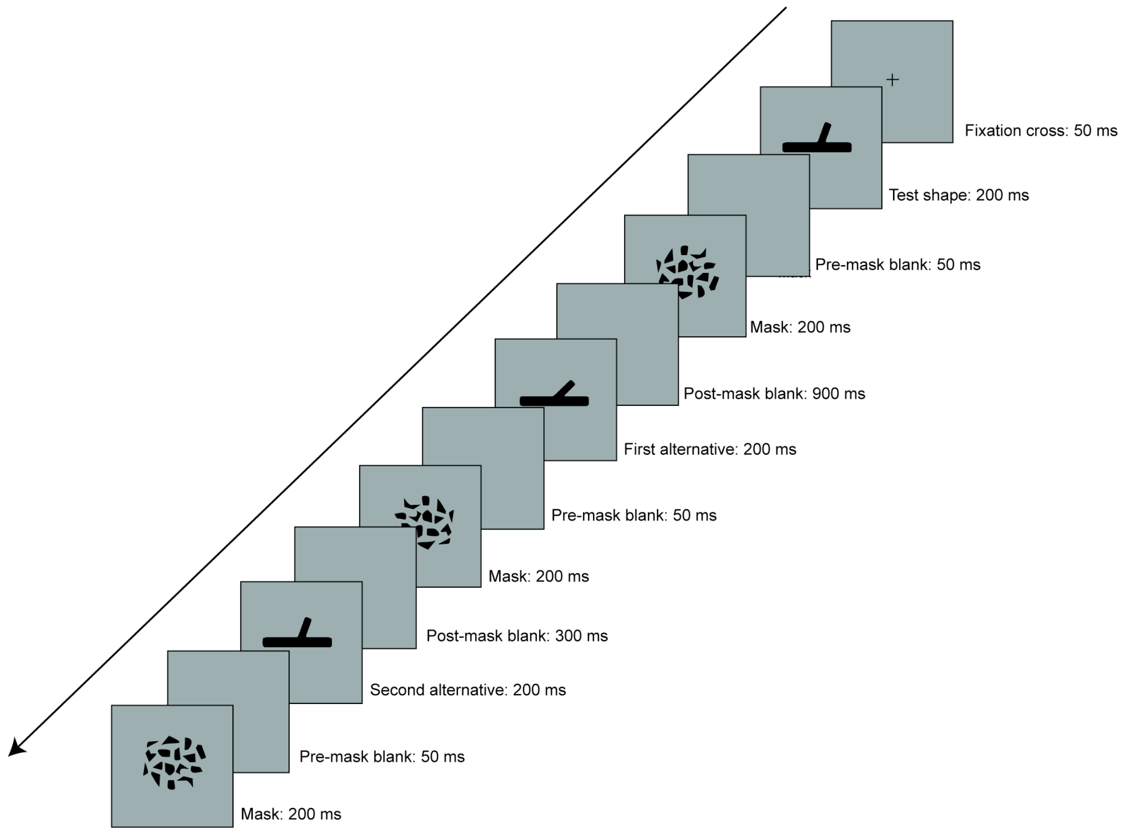

Figure 23. Trial sequence. Observers viewed a total of 3 shapes. The first shape was a "test" shape followed by 2 "alternative" shapes. The task of the observer was to choose which interval, first or second, contains the alternative that matches the test shape. All three shapes were followed by a mask.

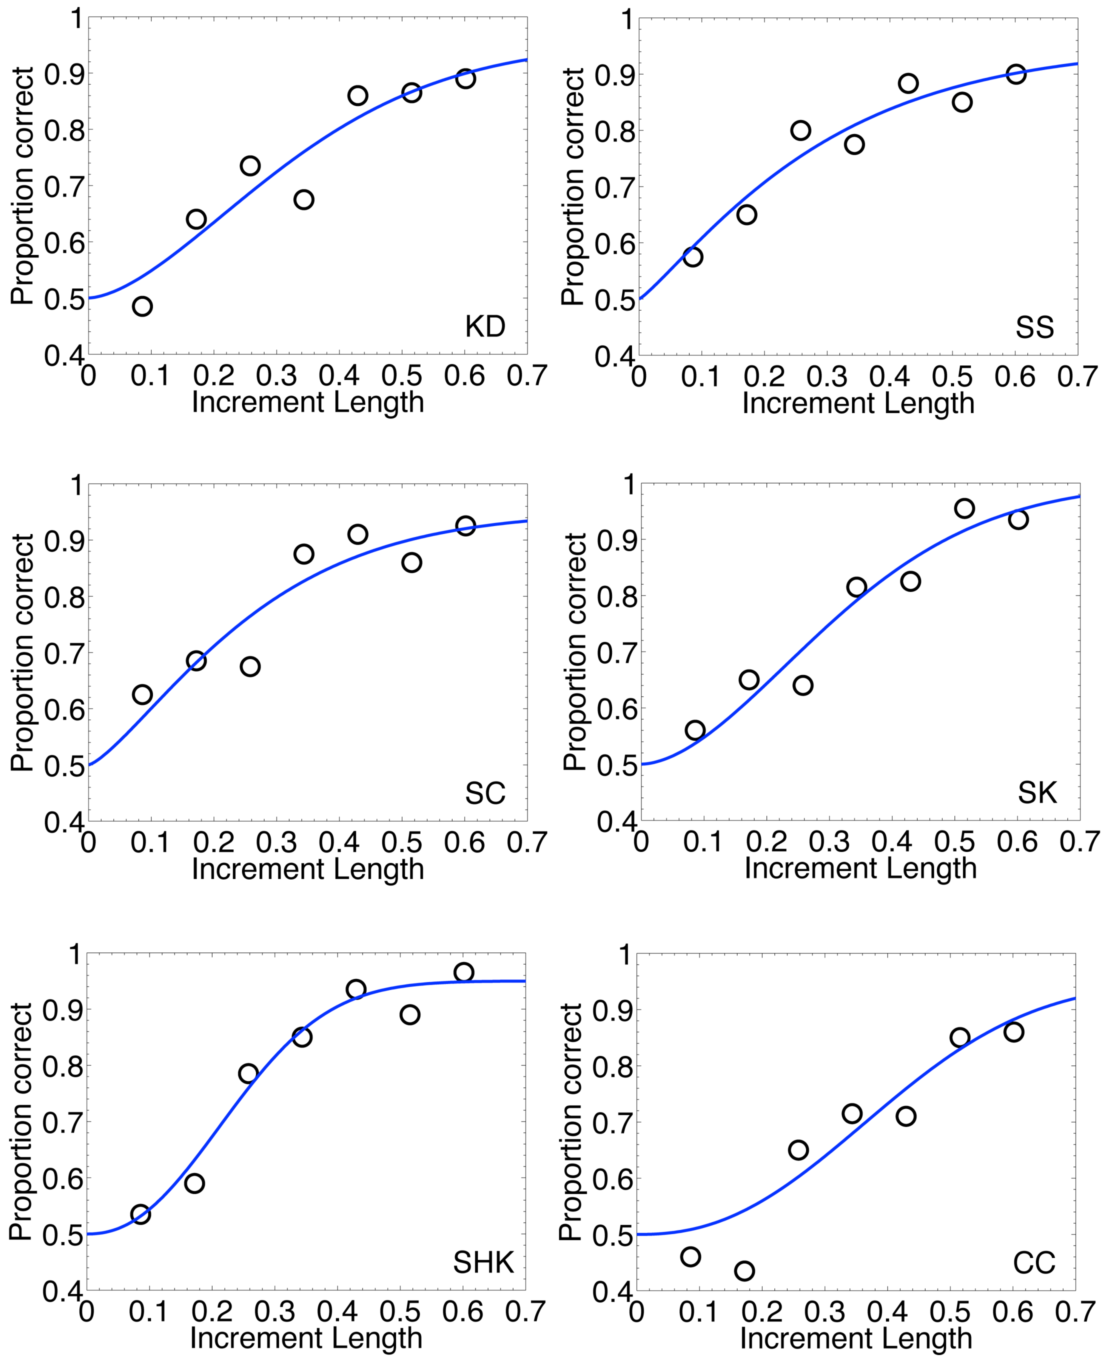

Figure 24a. Two-part shape (Experiment 1): Length transformation. Psychometric fits are shown for six observers.

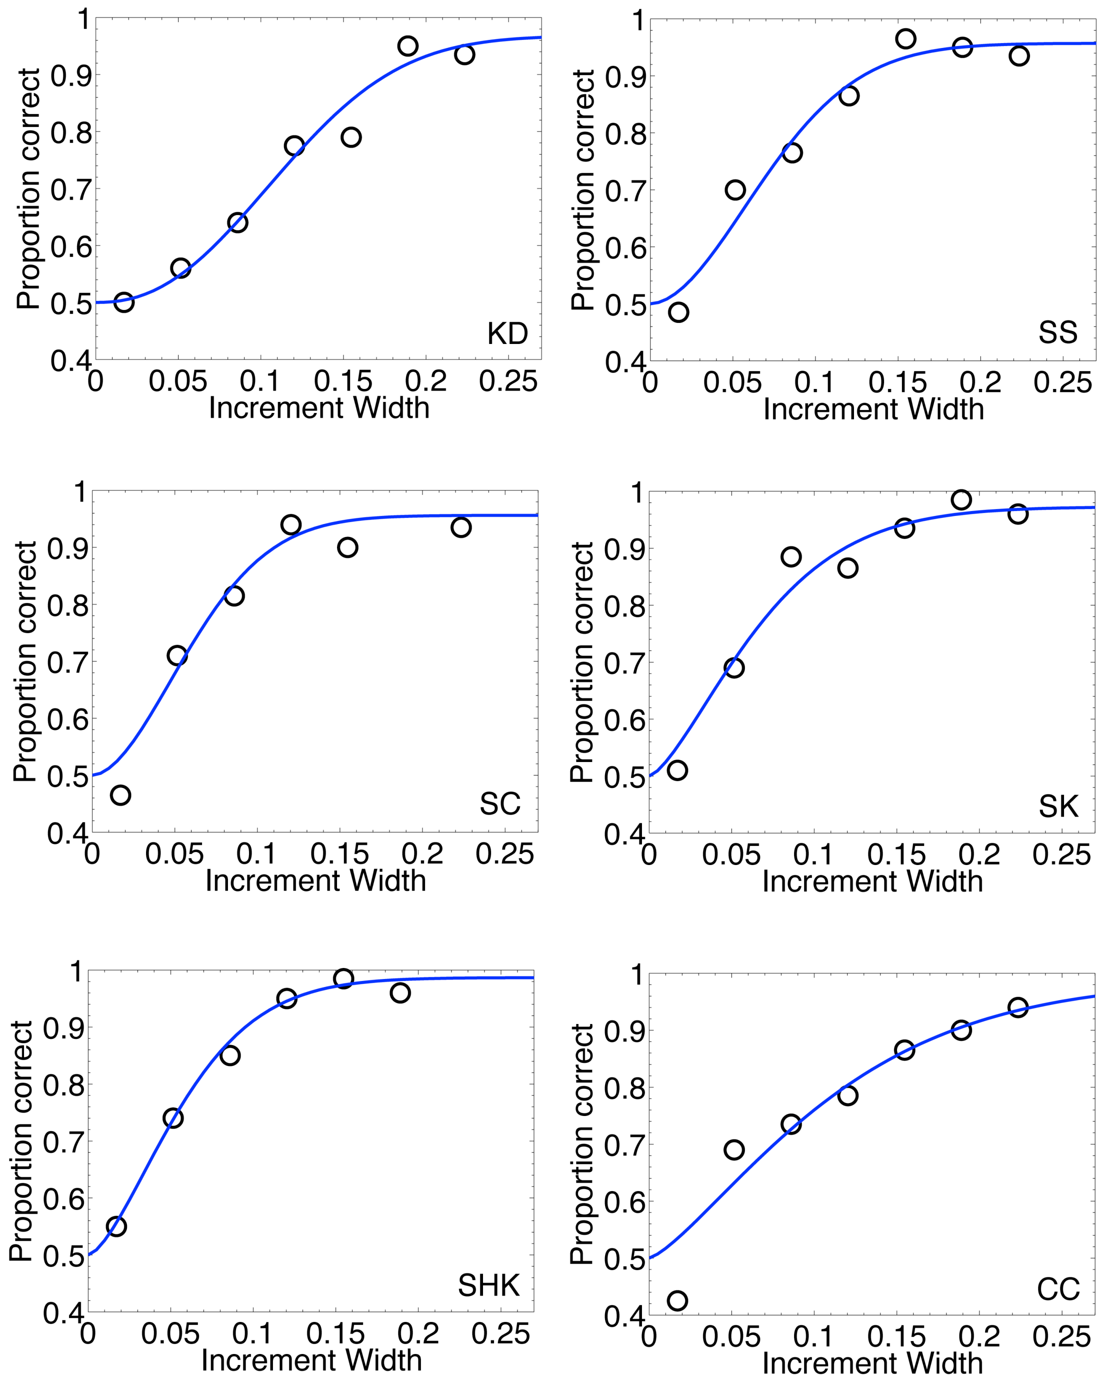

Figure 24b. Two-part shape (Experiment 1): Width transformation. Psychometric fits are shown for six observers.

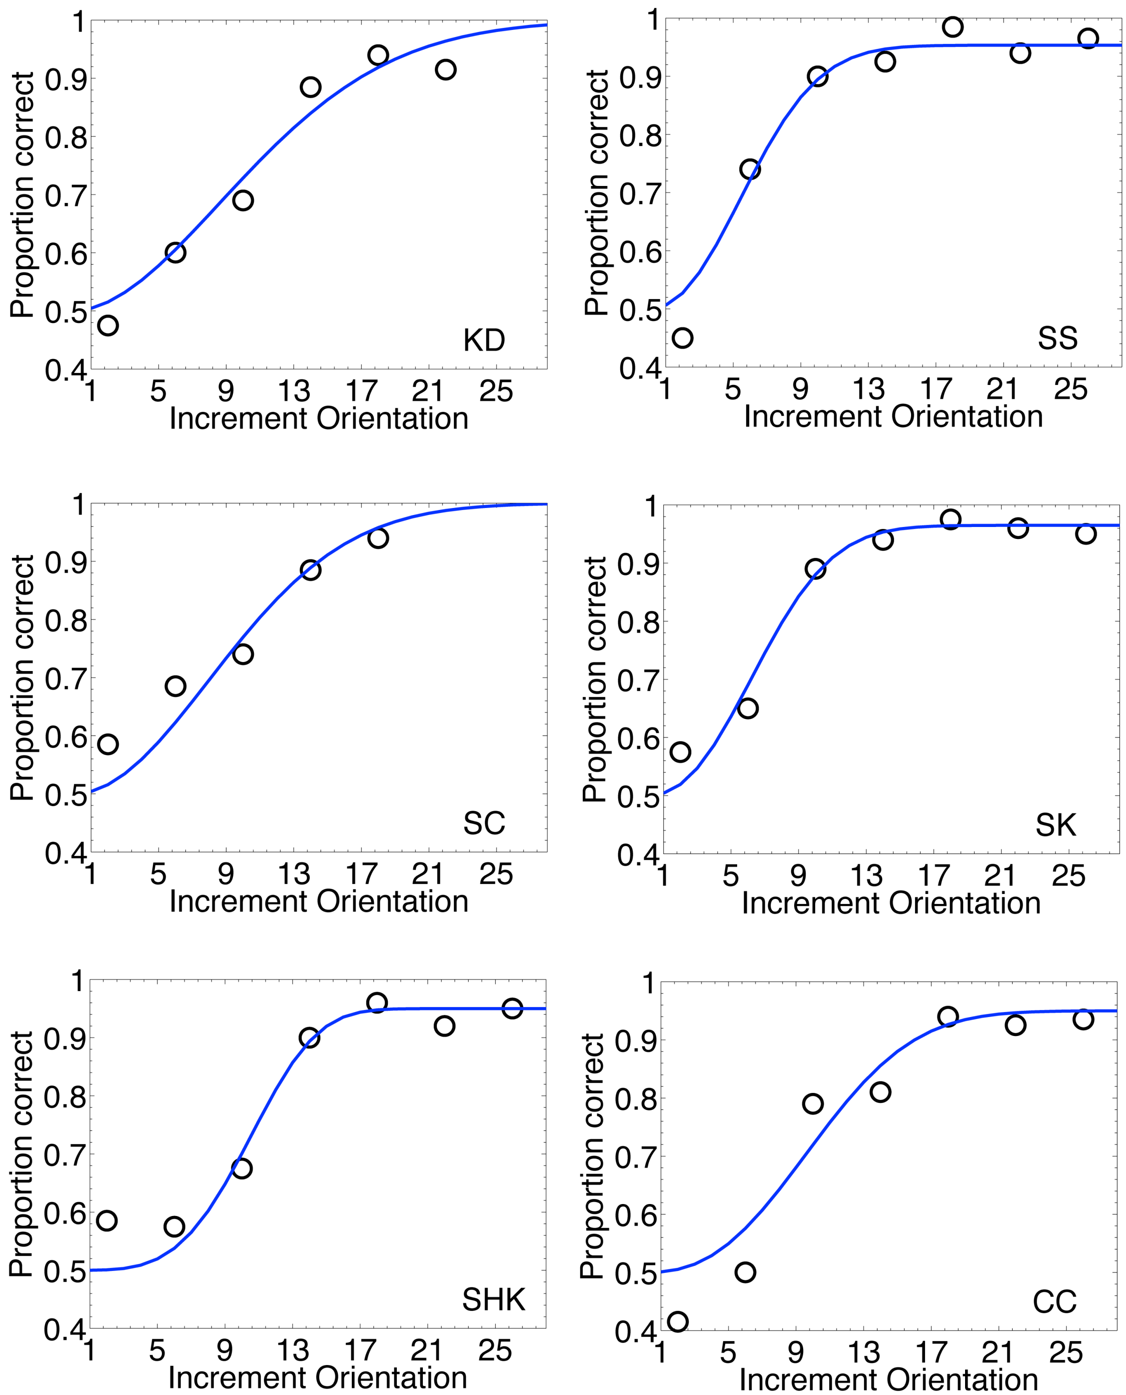

Figure 24c. Two-part shape (Experiment 1): Orientation transformation. Psychometric fits are shown for six observers.

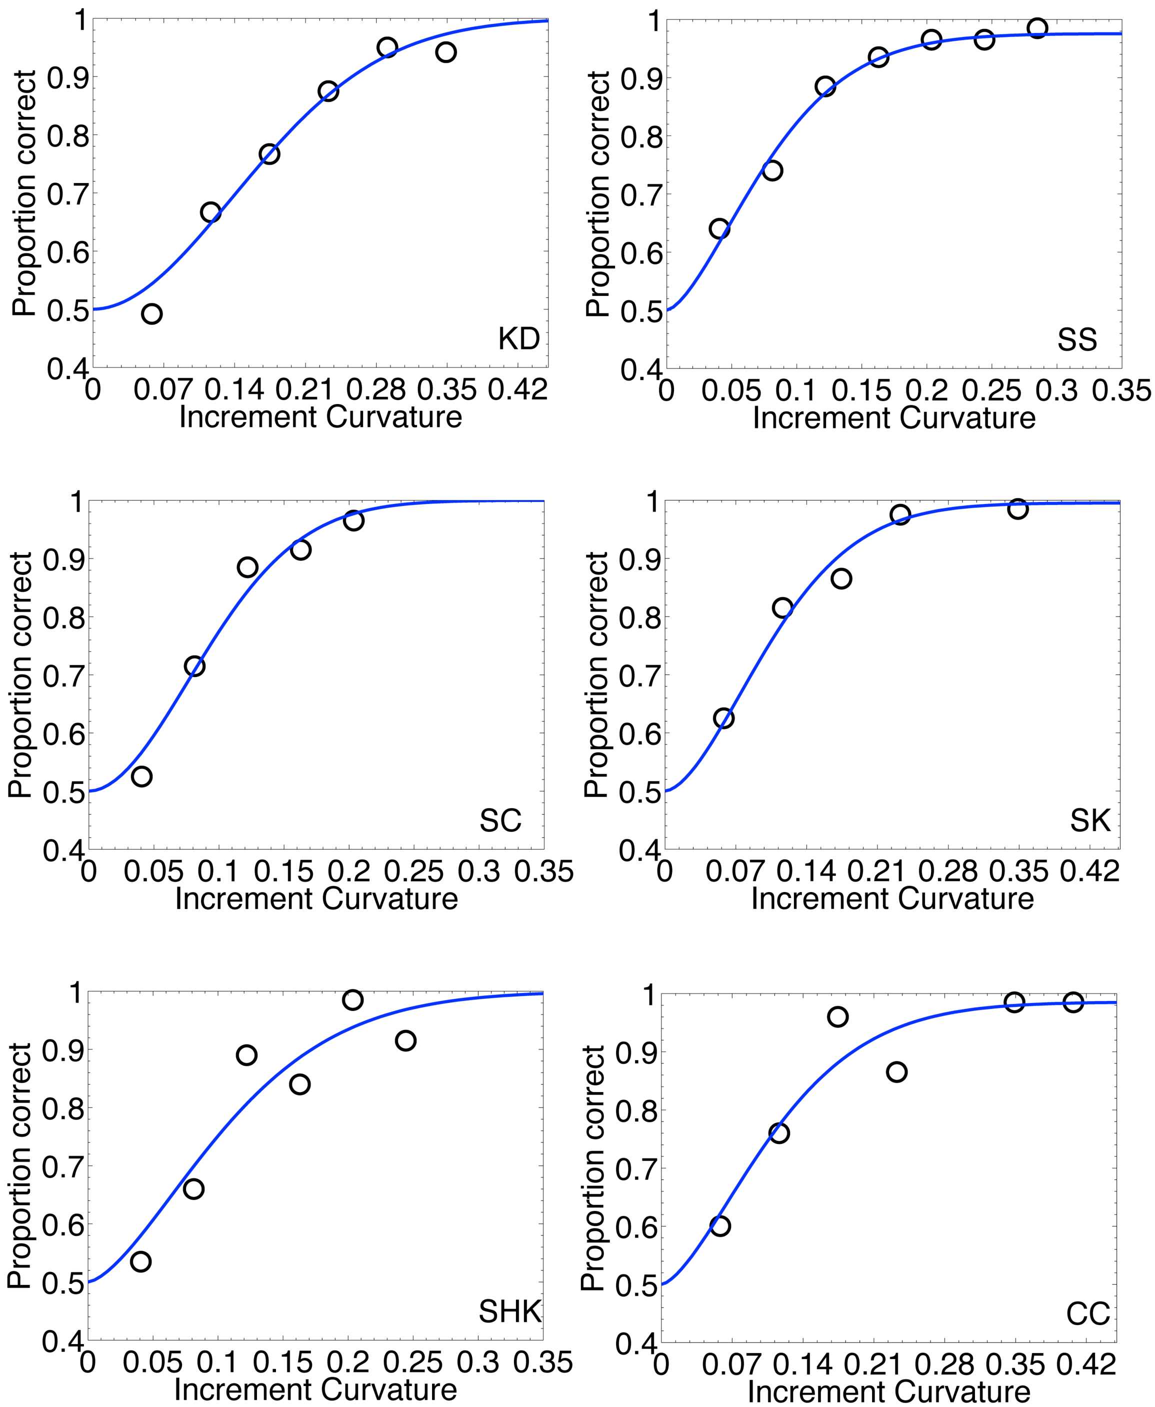

Figure 24d. Two-part shape (Experiment 1): Curvature transformation. Psychometric fits are shown for six observers.

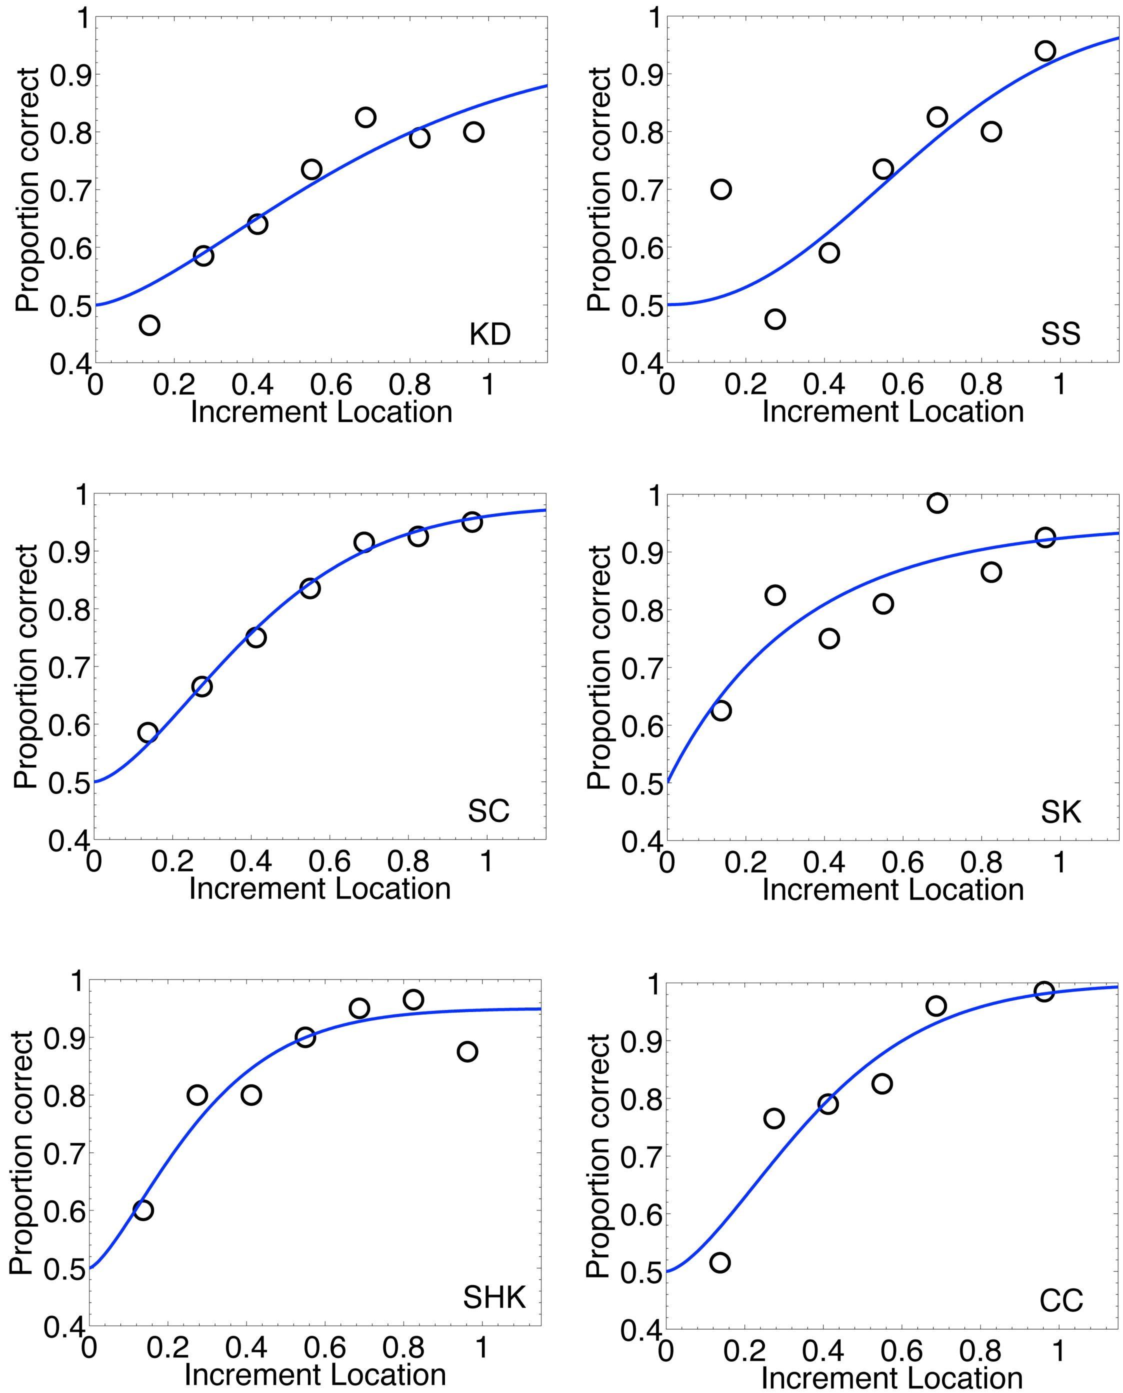

Figure 24e. Two-part shape (Experiment 1): Location transformation. Psychometric fits are shown for six observers.

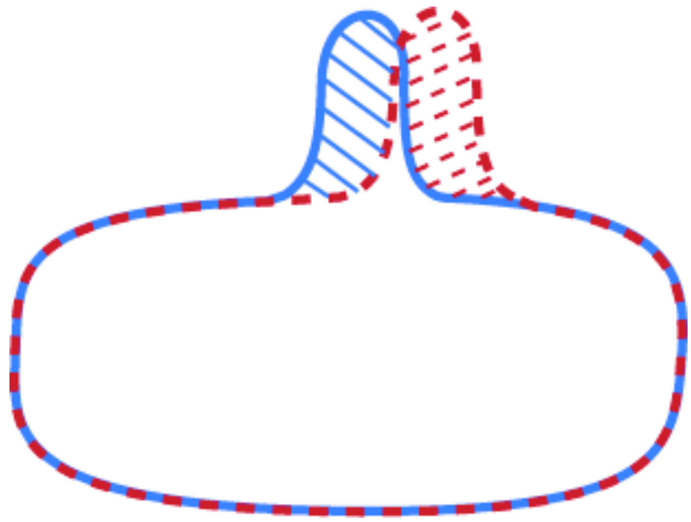

Figure 25. Illustration of an area-based shape difference measure for a two-part shape. The two shapes (Sh$_1$, blue shape and Sh$_2$, red shape: dotted line) are aligned maximally with one another; the non-overlapping regions are depicted by the left-facing (blue) and right-facing (red) hashed regions; those regions represent a difference in area between the two shapes. The sum of these two areas is then normalized by the sum of the areas of the two parts of the shapes to produce the shape difference metric, a value that illustrates the percent difference between the two shapes. The physical difference between the two shapes—the actual extent of the distance between the two parts is equal to the difference threshold for that condition. (Please refer to the formula in text).

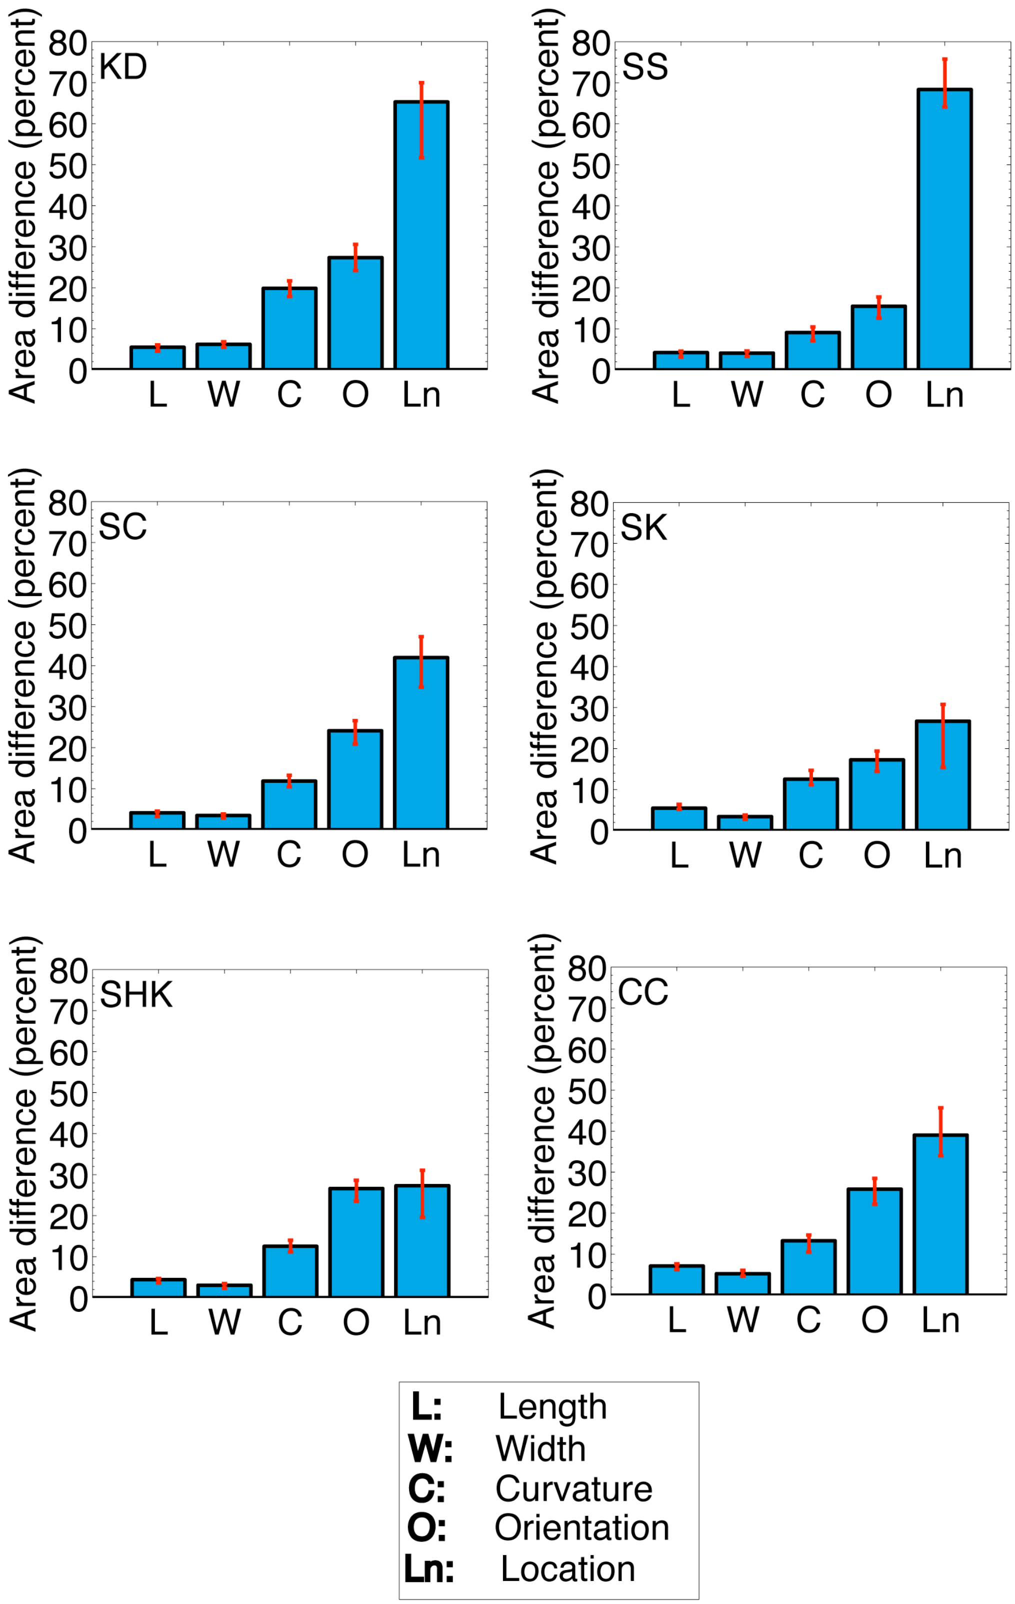

Figure 26. Experiment 1. Area difference plots for five conditions are shown for the six subjects. Observers are most sensitive to the width and length transformations, followed by curvature, and then by orientation. Observers are least sensitive to location transformation.

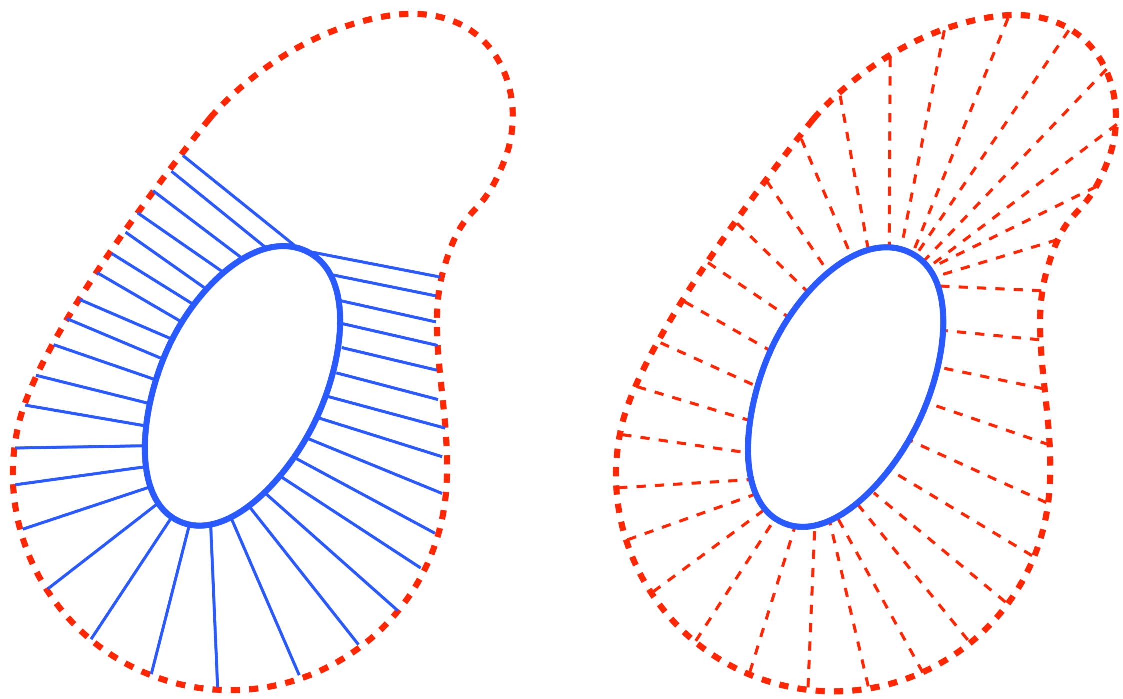

Figure 27. Illustration of the distance-based average common metric for two shapes. The two shapes (Sh_1, small blue shape) and (Sh_2, large red shape) are first aligned at their centers of mass. Then, the closest distance is computed from points on Sh_1 to corresponding points on Sh_2, and from Sh_2 to Sh_1. The average of these values is the final distance measure.

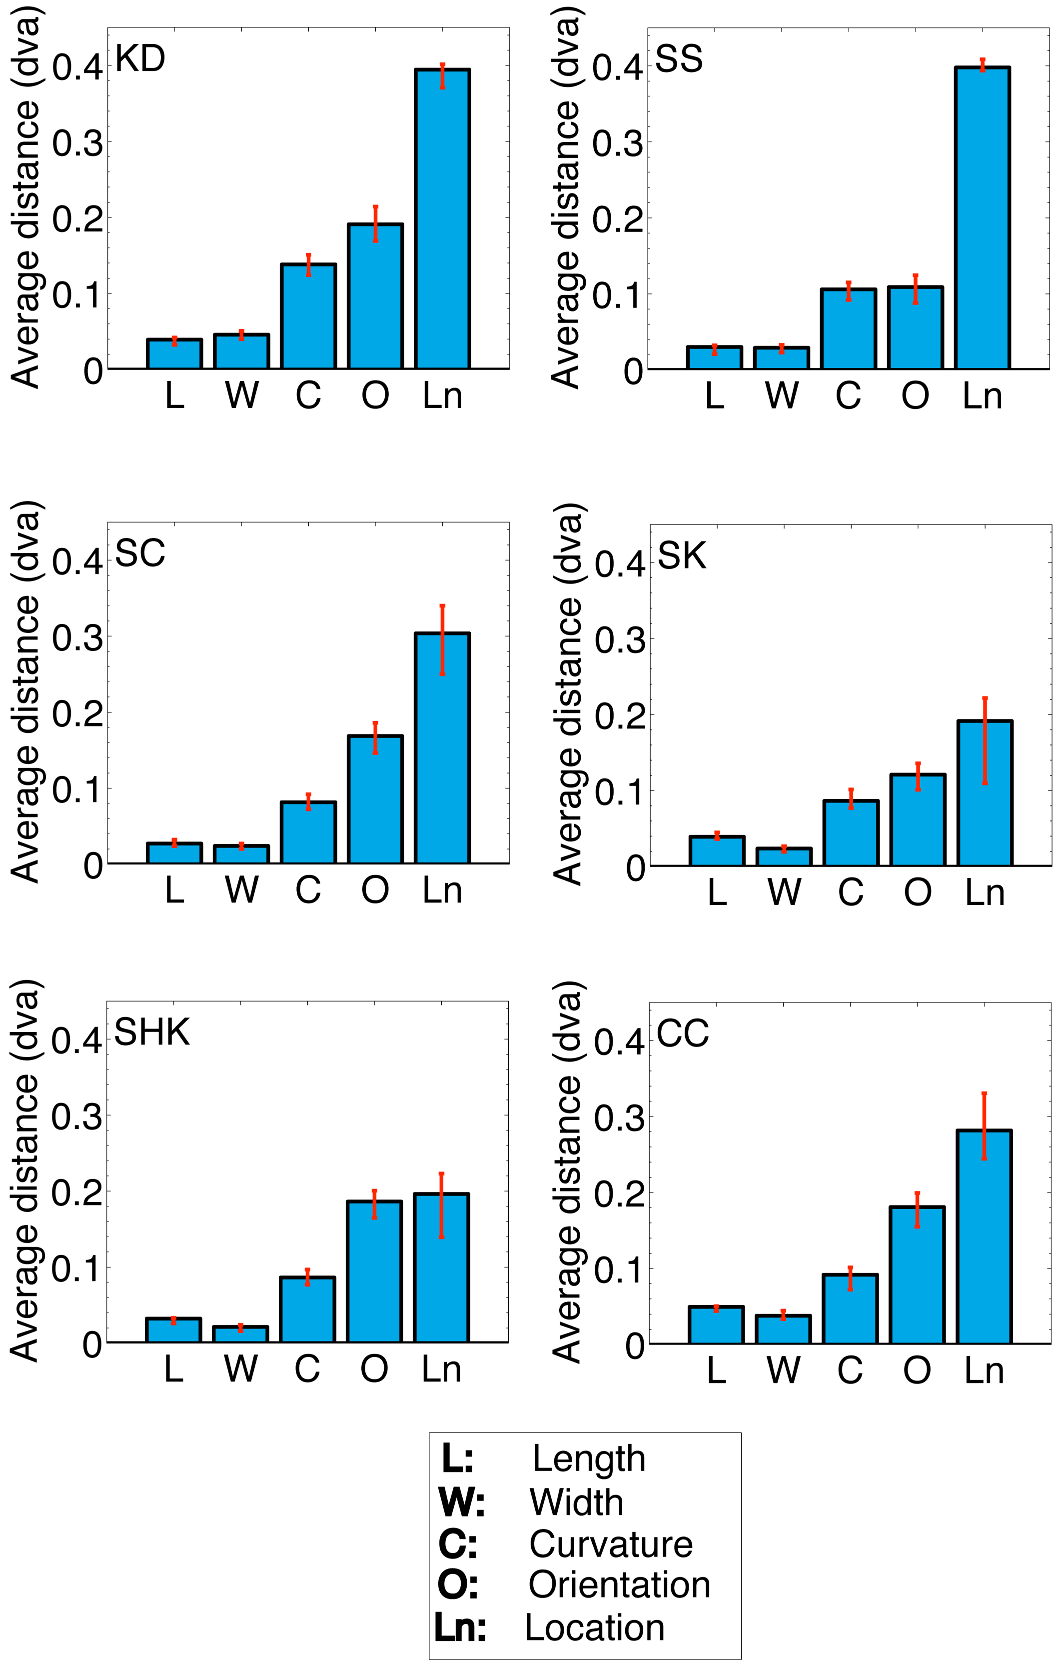

Figure 28. Experiment 1. Distance-based common metric plots for five conditions are shown for the six subjects. Observers are most sensitive to the width and length transformations, followed by curvature, and then by orientation. Observers are least sensitive to location transformation.

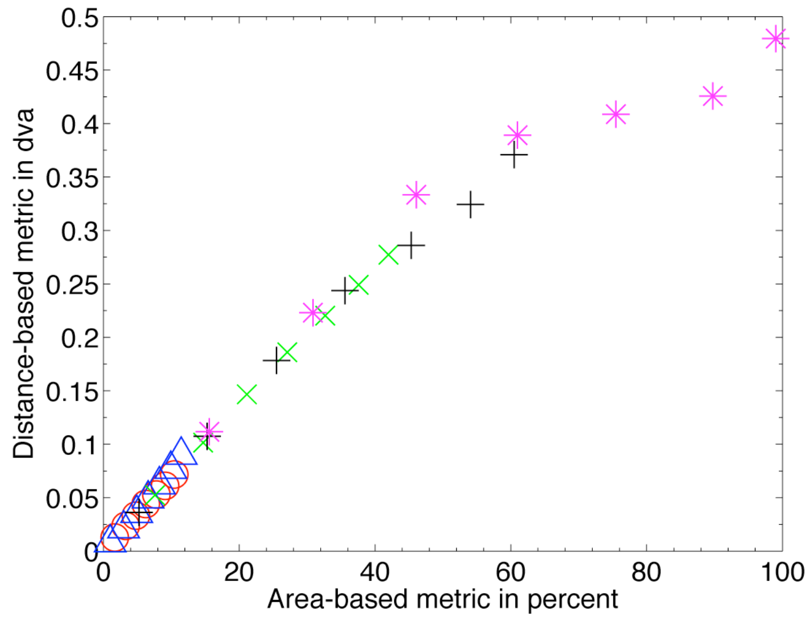

(a)

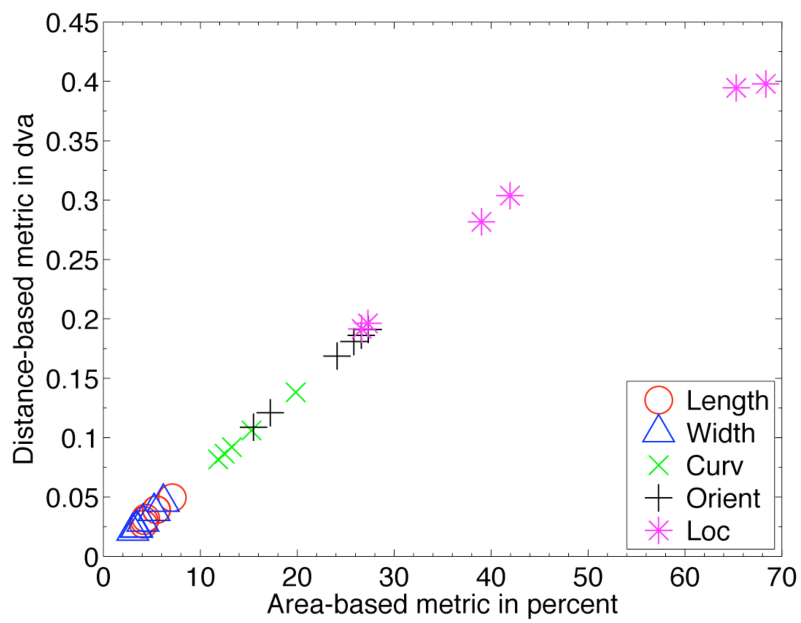

(b)

Figure 29. Correlation between area-based (x-axis) and distance-based (y-axis) common metrics for length, width, curvature, orientation, and location transformations. (a): Correlation between the increment values converted into area-based and distance-based metrics; a total of 7 increment values for each transformation was used. (b): Correlation between the thresholds converted into area-based and distance-based for 6 observers.

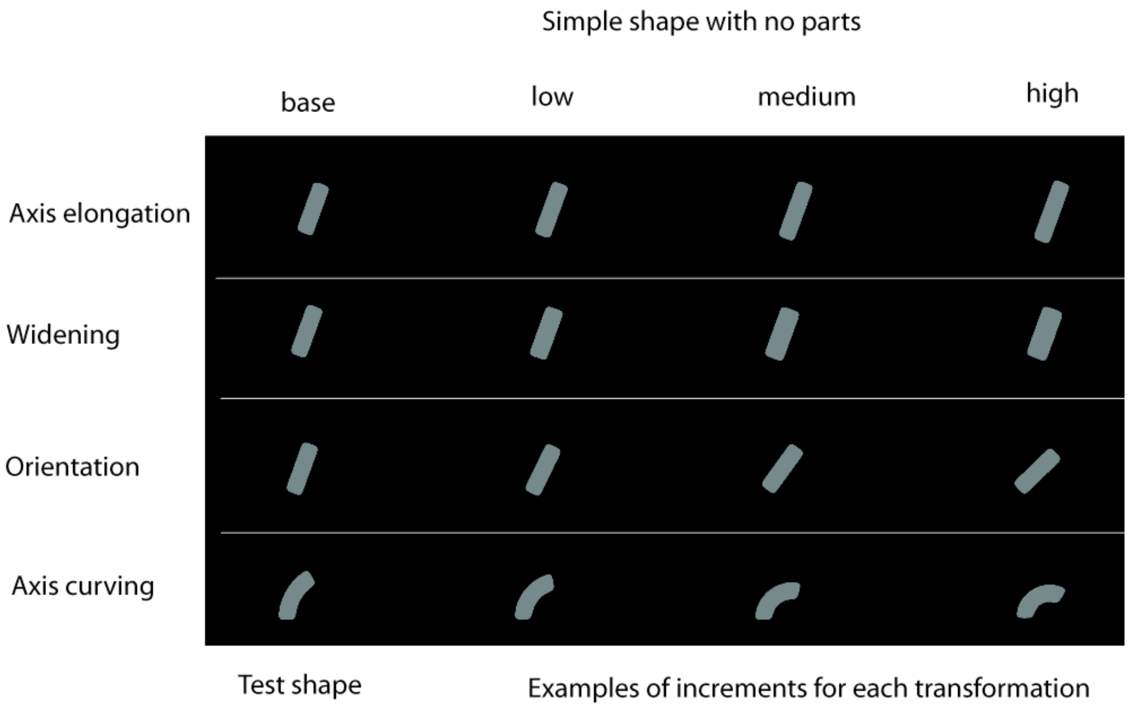

Figure 30. Experiment 2: four transformations are shown for the simple, single-axis shape: axis elongation (length), width, orientation, and curvature (axis curving).

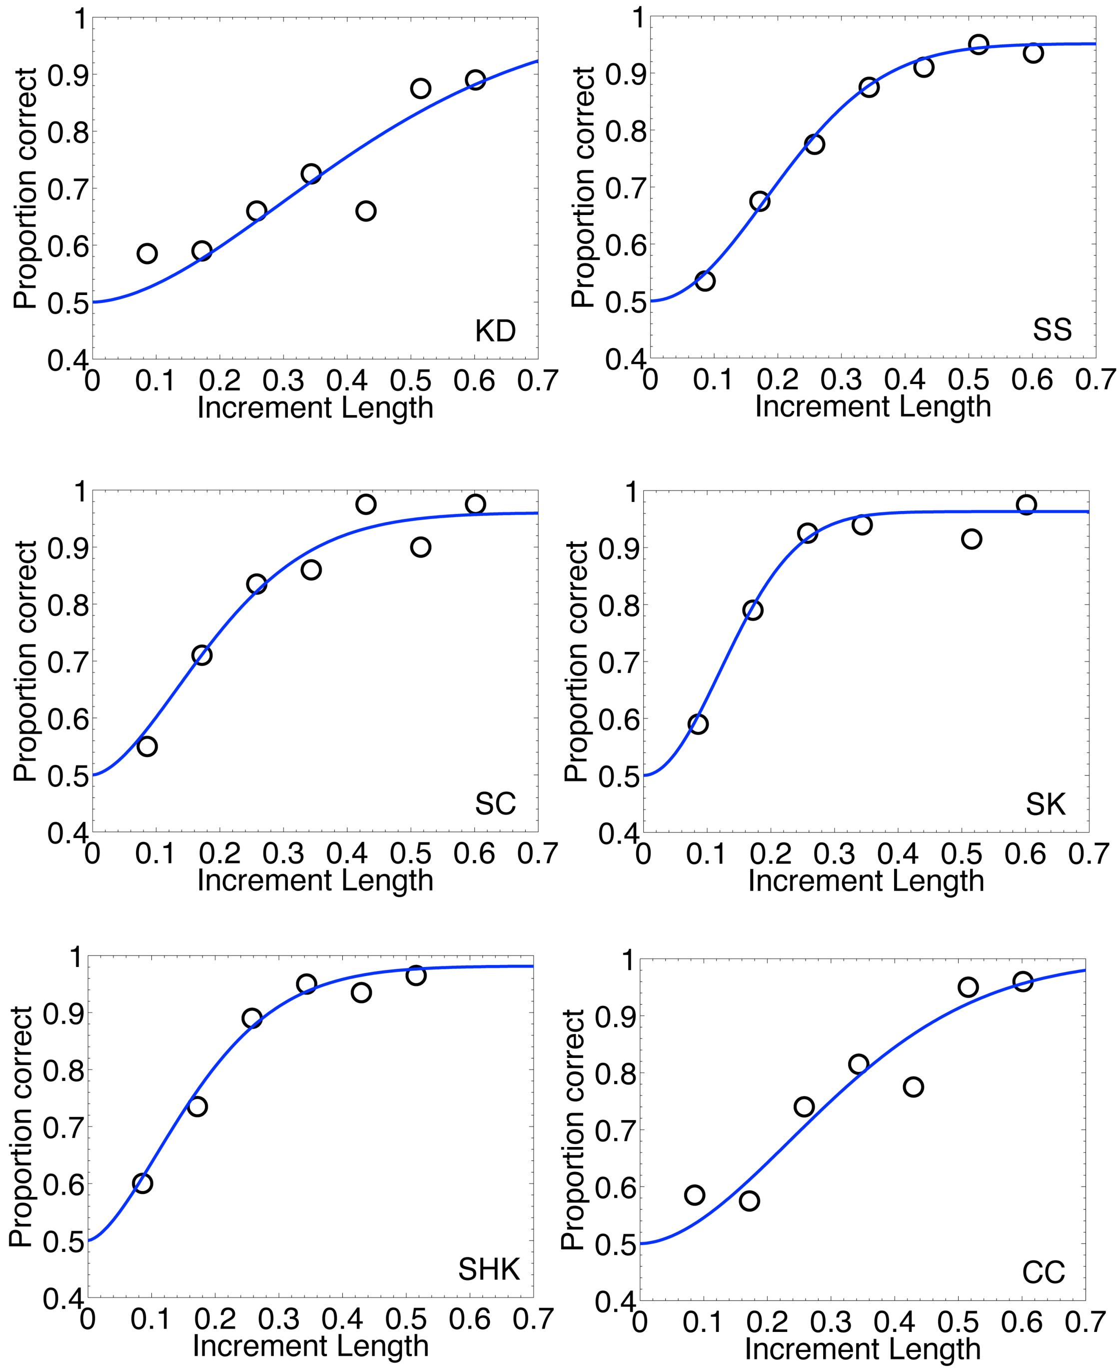

Figure 31a. Single-axis shape: Length transformation. Psychometric fits are shown for six observers.

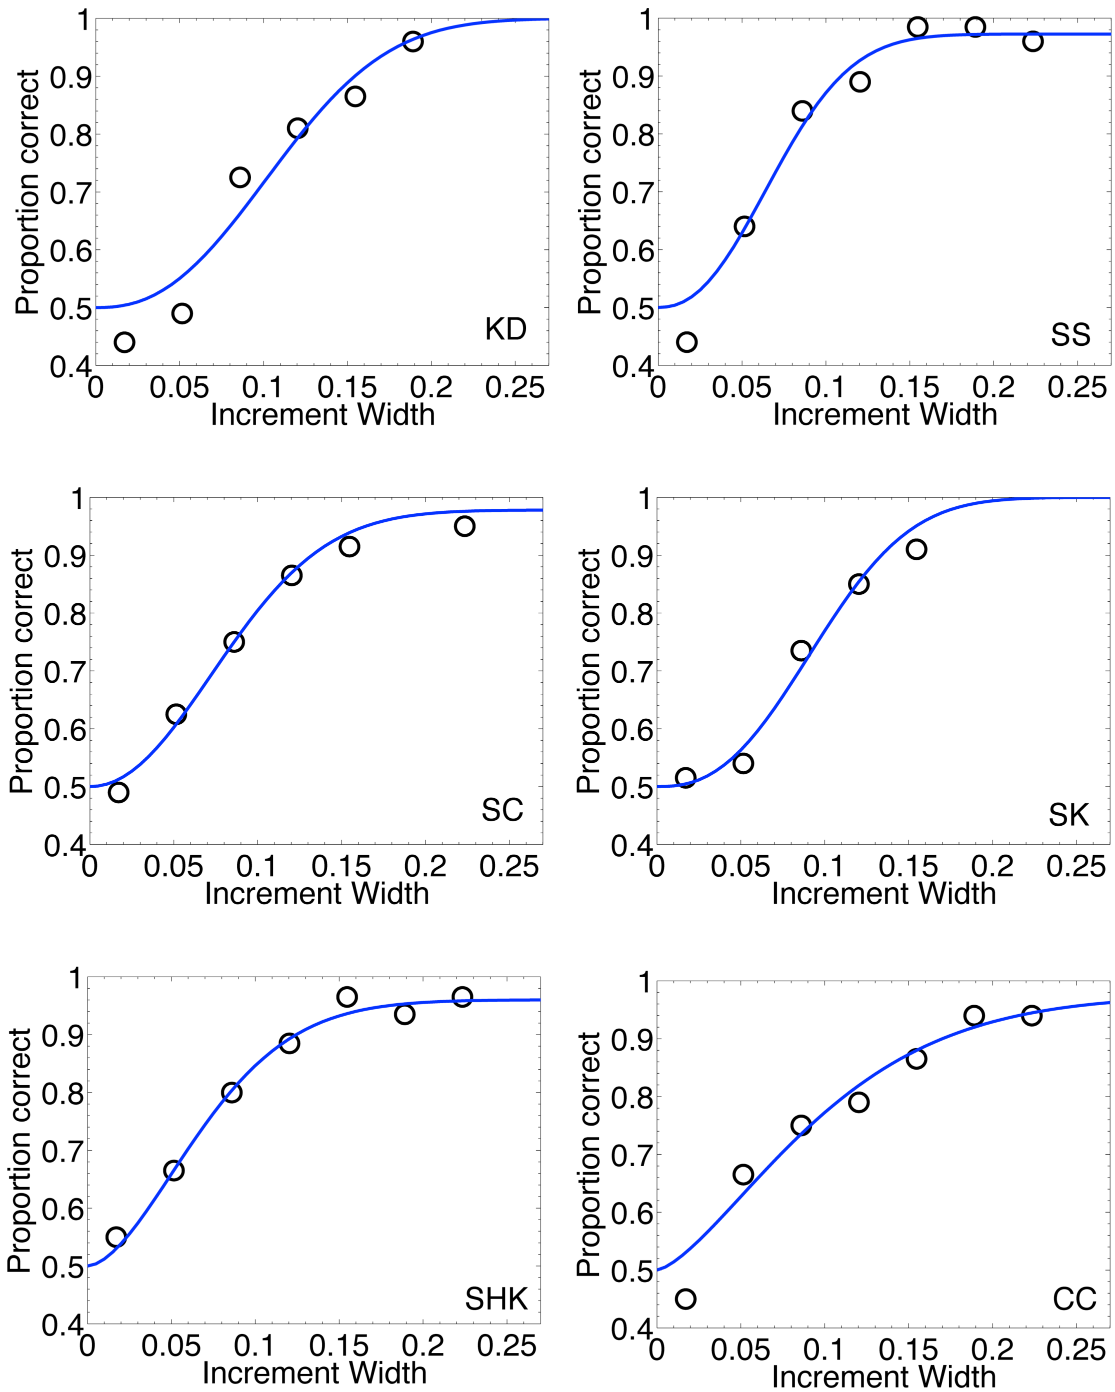

Figure 31b. Single-axis shape: Width transformation. Psychometric fits are shown for six observers.

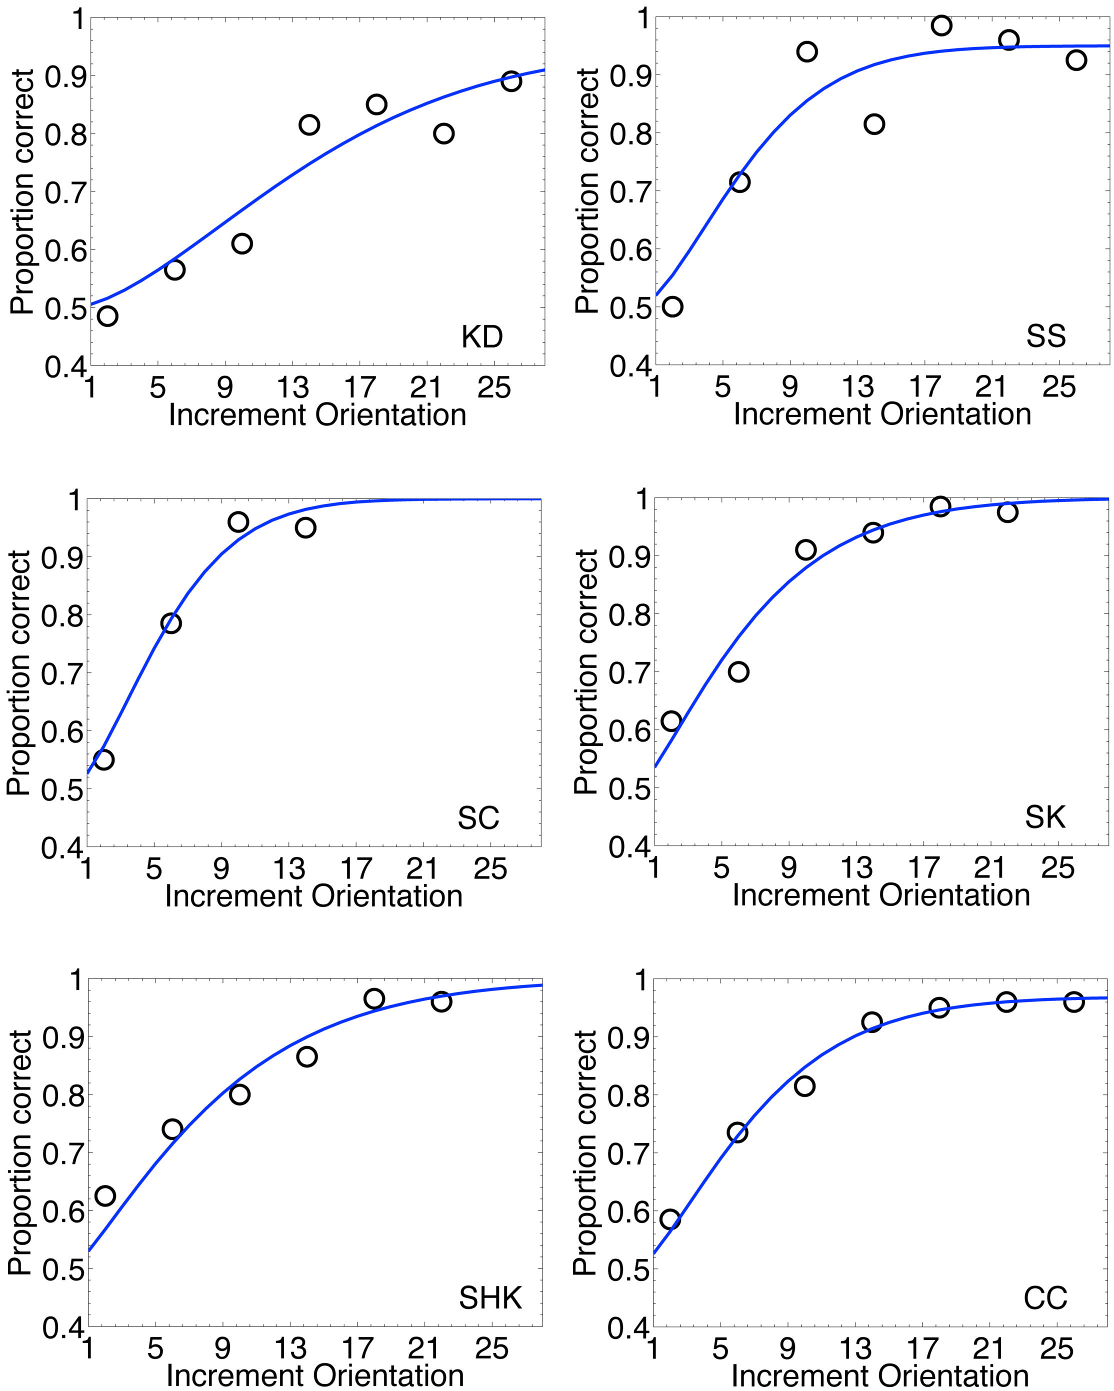

Figure 31c. Single-axis shape: Orientation transformation. Psychometric fits are shown for six observers.

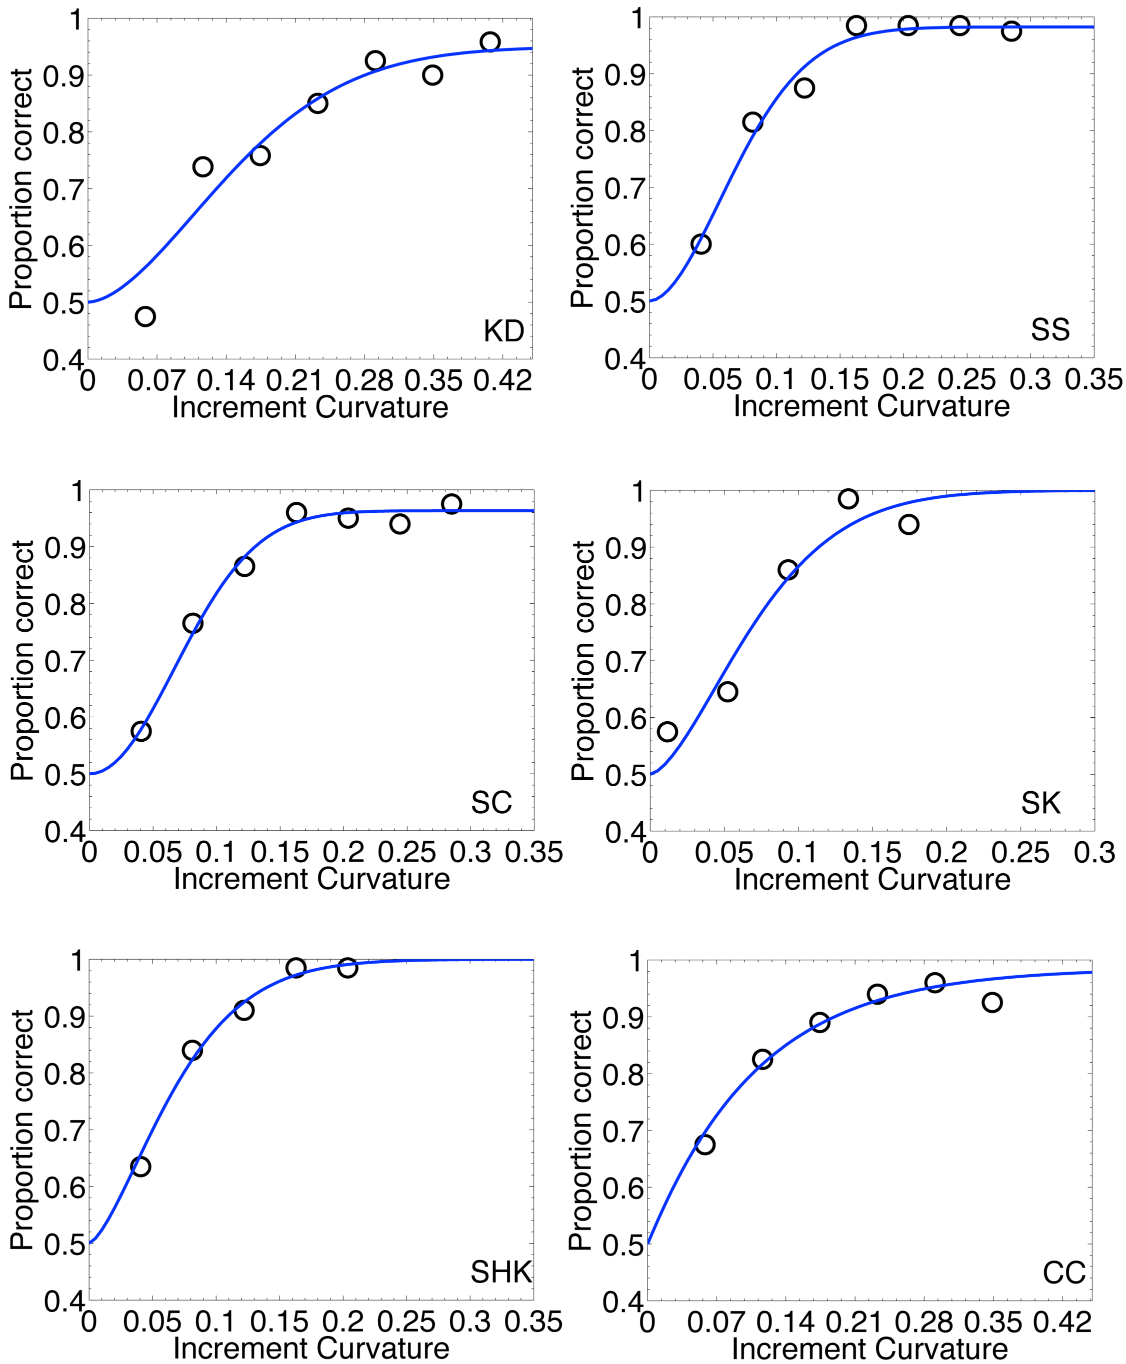

Figure 31d. Single-axis shape: Curvature transformation. Psychometric fits are shown for six observers.

Figure 32. Schematic illustration of an area-based shape difference measure for a single-axis shape (length transformation example). This procedure is same as that used in Experiment 1. The two shapes (Sh_1, designated by the black continuous line and Sh_2, black dotted line) are aligned along their bottom portion, similar to the alignment of the parts in the two-part shape in Experiment 1. The non-overlapping regions are depicted by the left-facing and right-facing hashed regions; those regions represent a difference in area between the two shapes. The sum of these two areas is then normalized by the sum of the areas of the two shapes to produce the shape difference metric, a value that illustrates the percent difference between the two shapes.

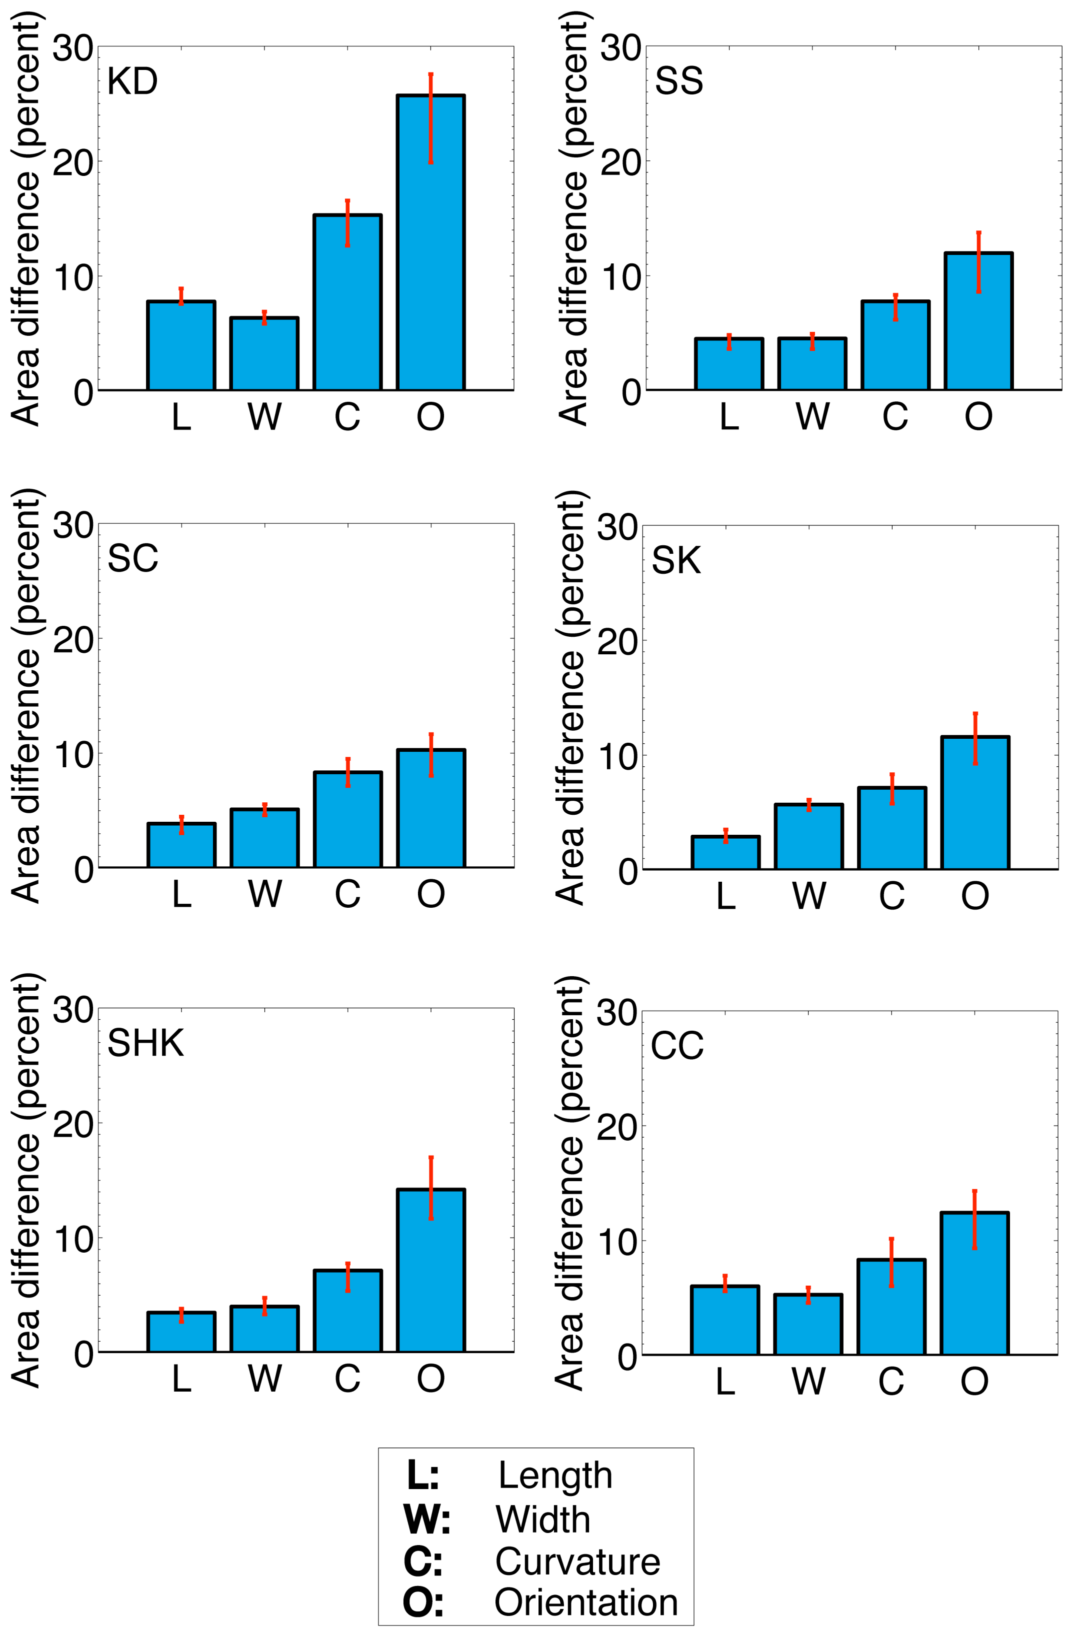

Figure 33. Experiment 2: Area difference plots for all four conditions are shown for the six subjects. Observers are least sensitive to the orientation transformation, more sensitive to the curvature transformation, and the most sensitive to the width and length transformation.

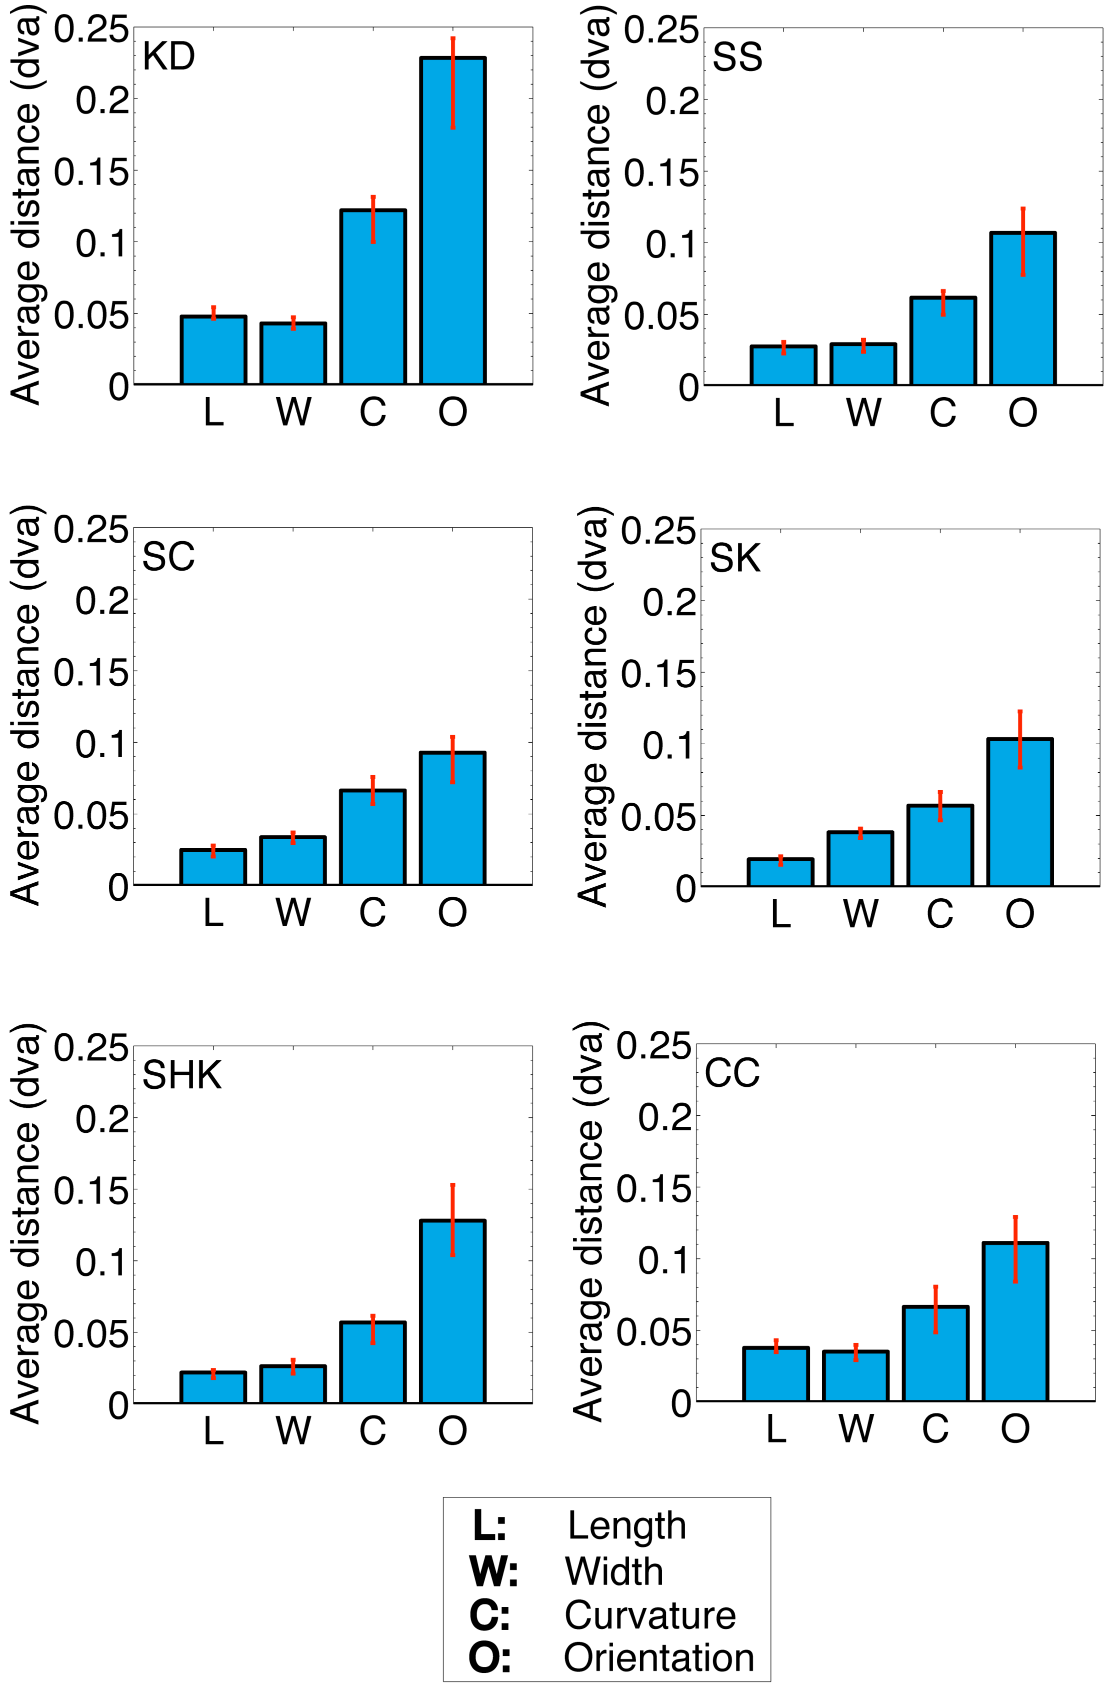

Figure 34. Experiment 2: Distance-based common metric plots for all four conditions are shown for the six subjects. Observers are least sensitive to the orientation transformation, more sensitive to the curvature transformation, and the most sensitive to the width and length transformation.

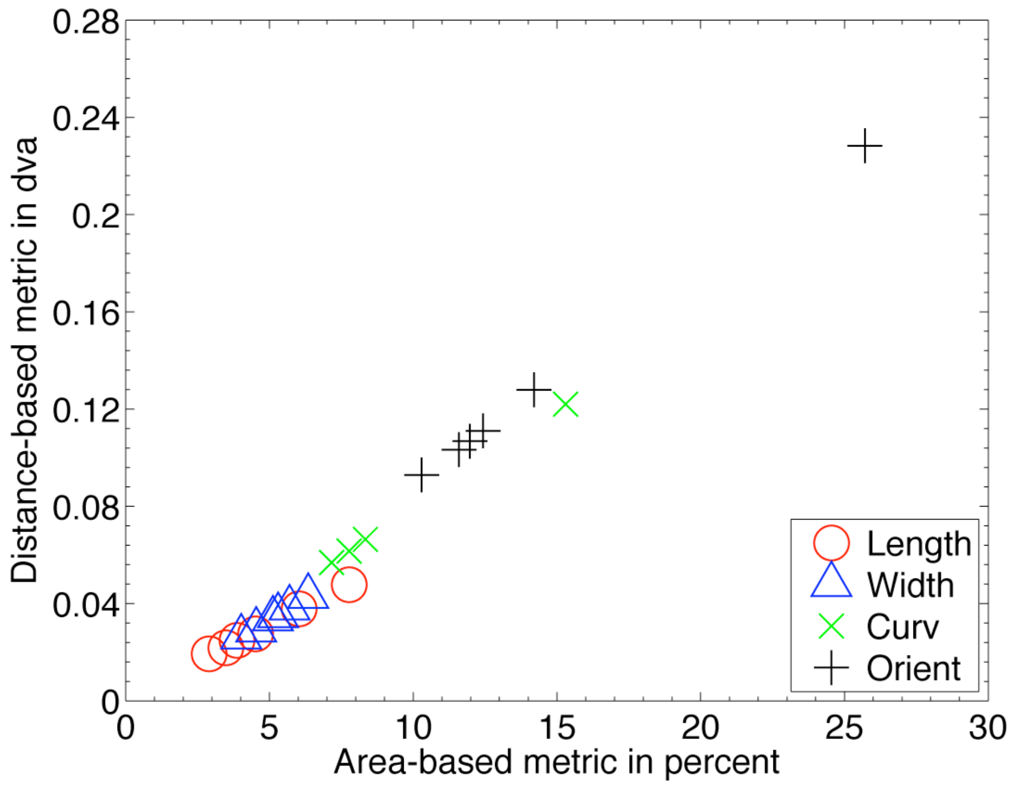

Figure 35. Correlation between area-based (x-axis) and distance-based (y-axis) common metrics for length, width, curvature, and orientation transformations: Experiment 2. The plot shows a total of 7 converted increment values for each metric, for 6 observers.

Figure 36. Schematic illustration of an area-based shape difference measure for a single-axis shape when the two shapes are maximally aligned with one another (length transformation example). First, the two shapes (Sh_1, designated by the black continuous line and Sh_2, black dotted line) are maximally aligned by their centers of mass. Next, the procedure is identical to the previous version (where shapes were not maximally aligned): the non-overlapping regions are depicted by the left-facing and right-facing hashed regions; those regions represent a difference in area between the two shapes. The sum of these two areas is then normalized by the sum of the areas of the two shapes to produce the shape difference metric, a value that illustrates the percent difference between the two shapes.

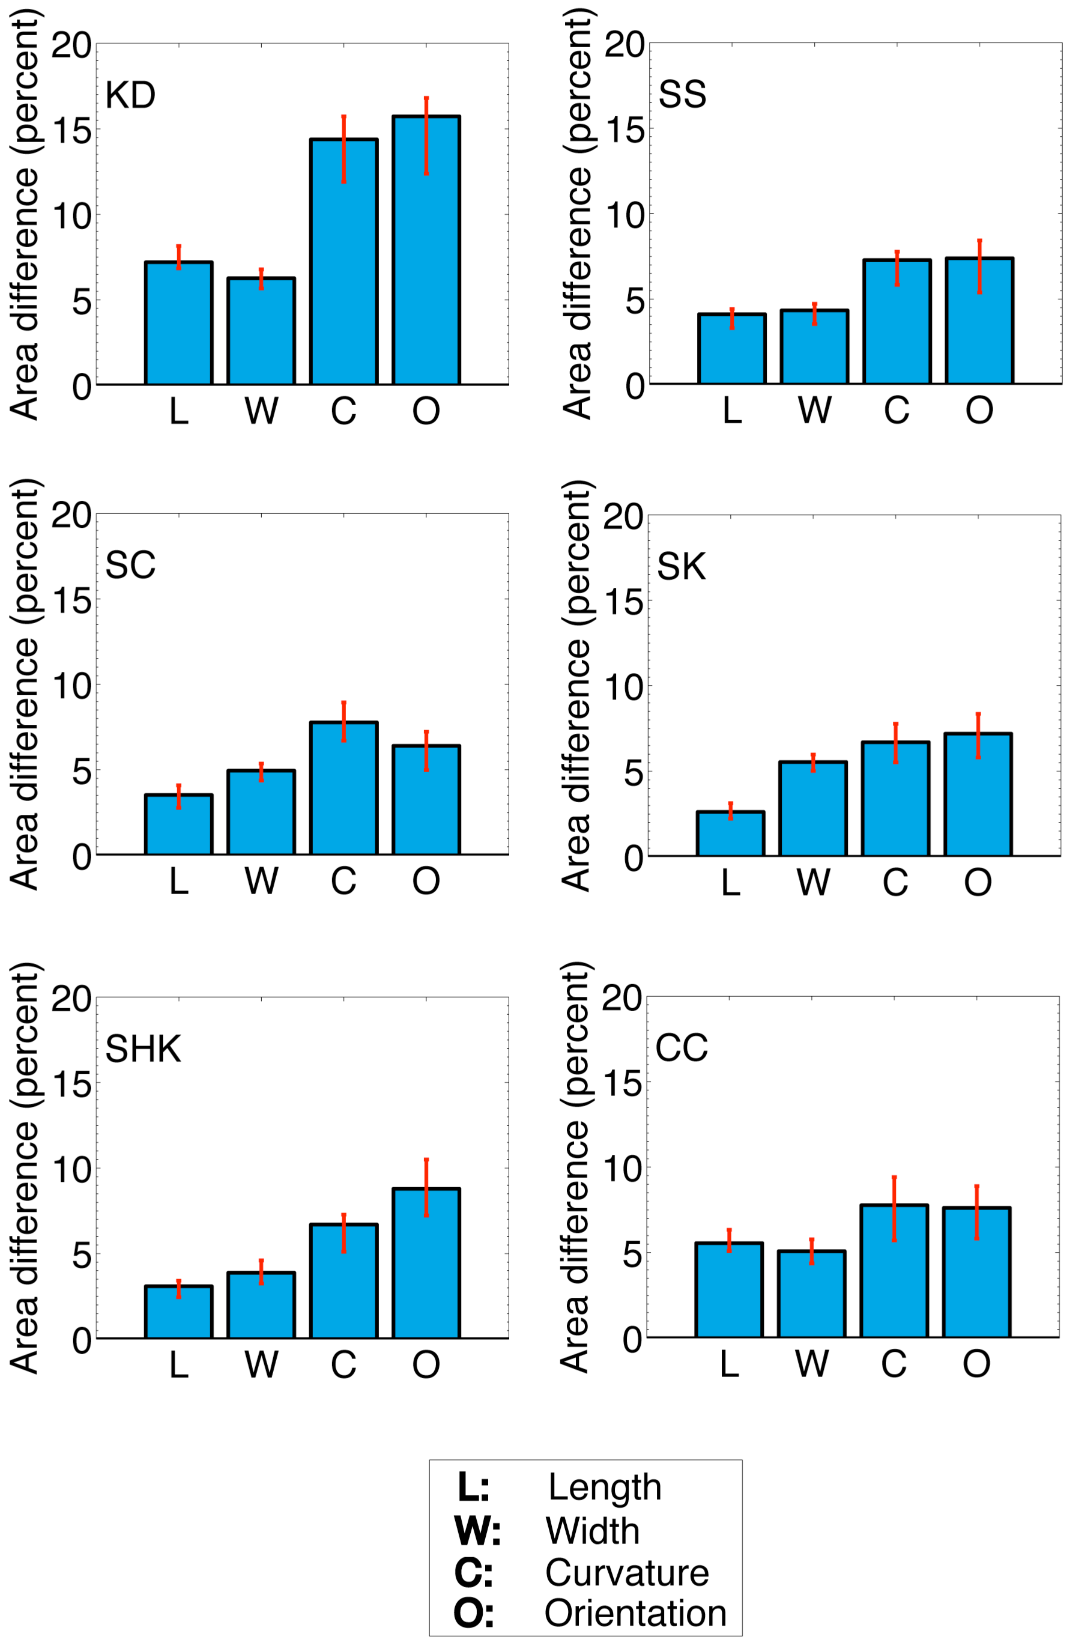

Figure 37. Experiment 2: Area difference plots for all four conditions are shown for the six observers; shapes were maximally aligned. Observers are most sensitive to width and length transformations, and are least sensitive to the orientation and curvature transformation.

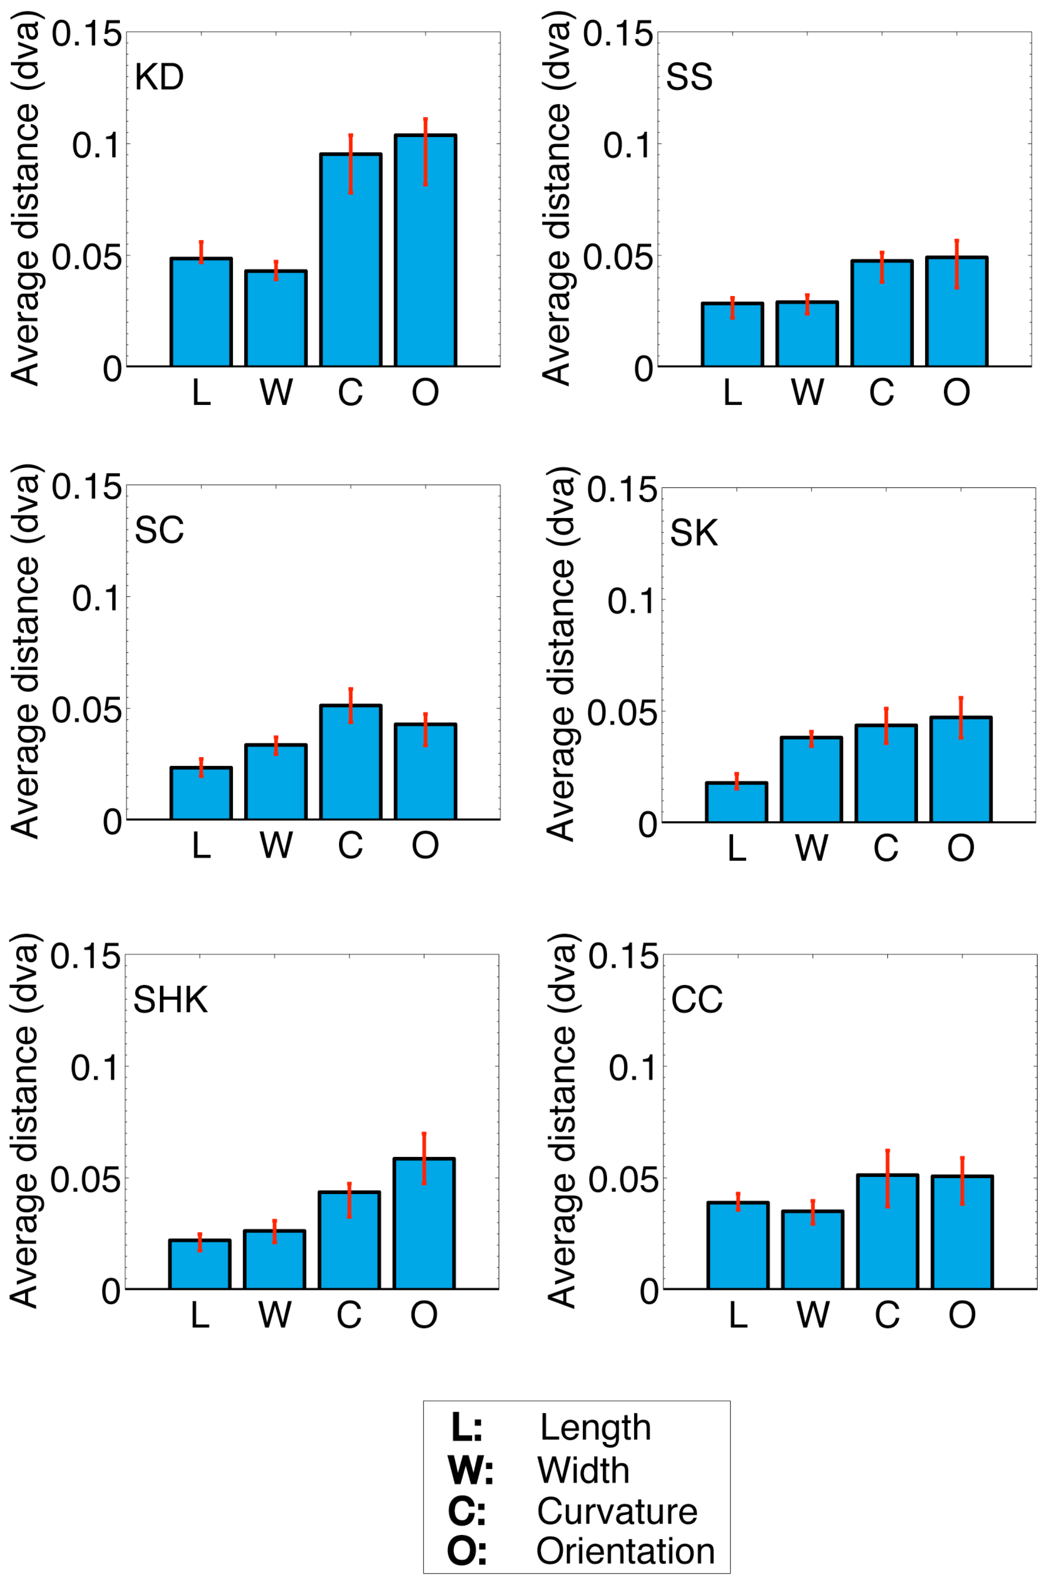

Figure 38. Experiment 2: Distance-based metric plots for all four conditions are shown for the six observers; shapes were maximally aligned. Observers are most sensitive to width and length transformations, and are least sensitive to the orientation and curvature transformation.

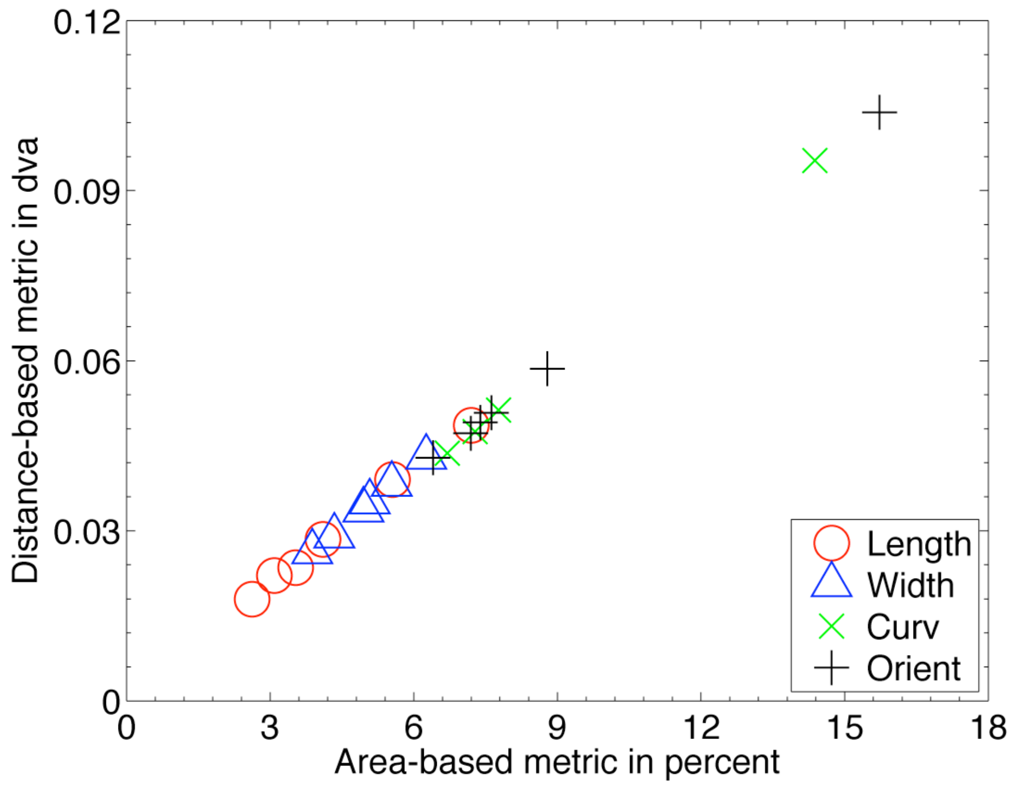

Figure 39. Correlation between area-based (x-axis) and distance-based (y-axis) common metrics for length, width, curvature, and orientation transformations: Experiment 2, maximally aligned shapes. The plot shows a total of 7 converted increment values for each metric, for 6 observers.

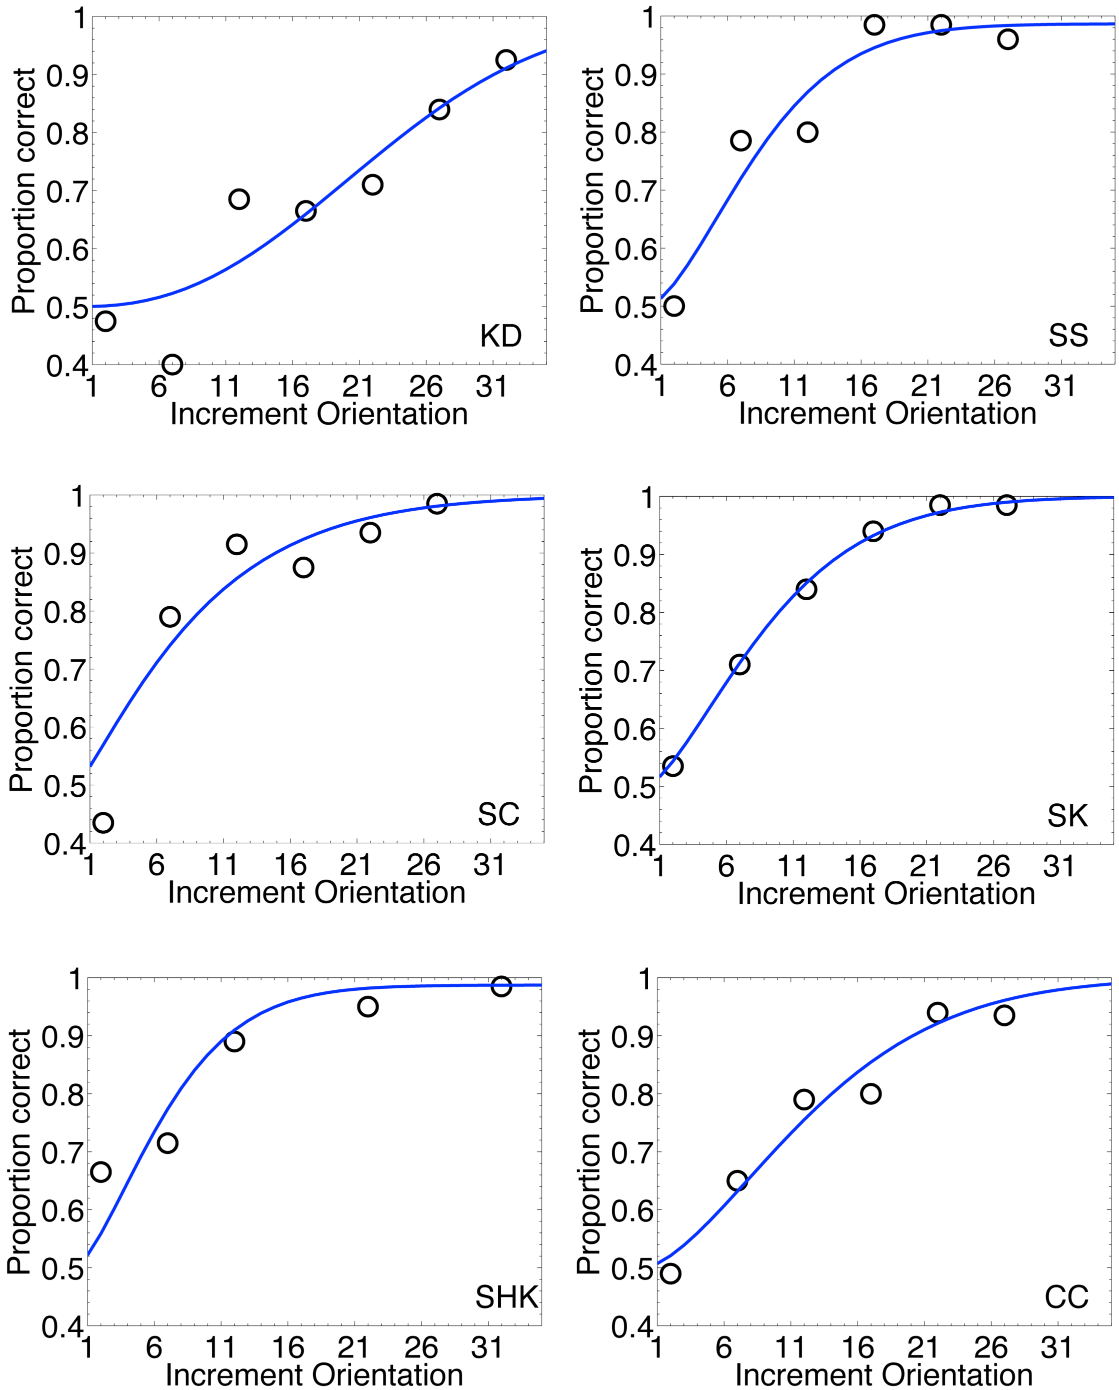

Figure 40a. Single-axis shape (Experiment 3a): Random orientation transformation.

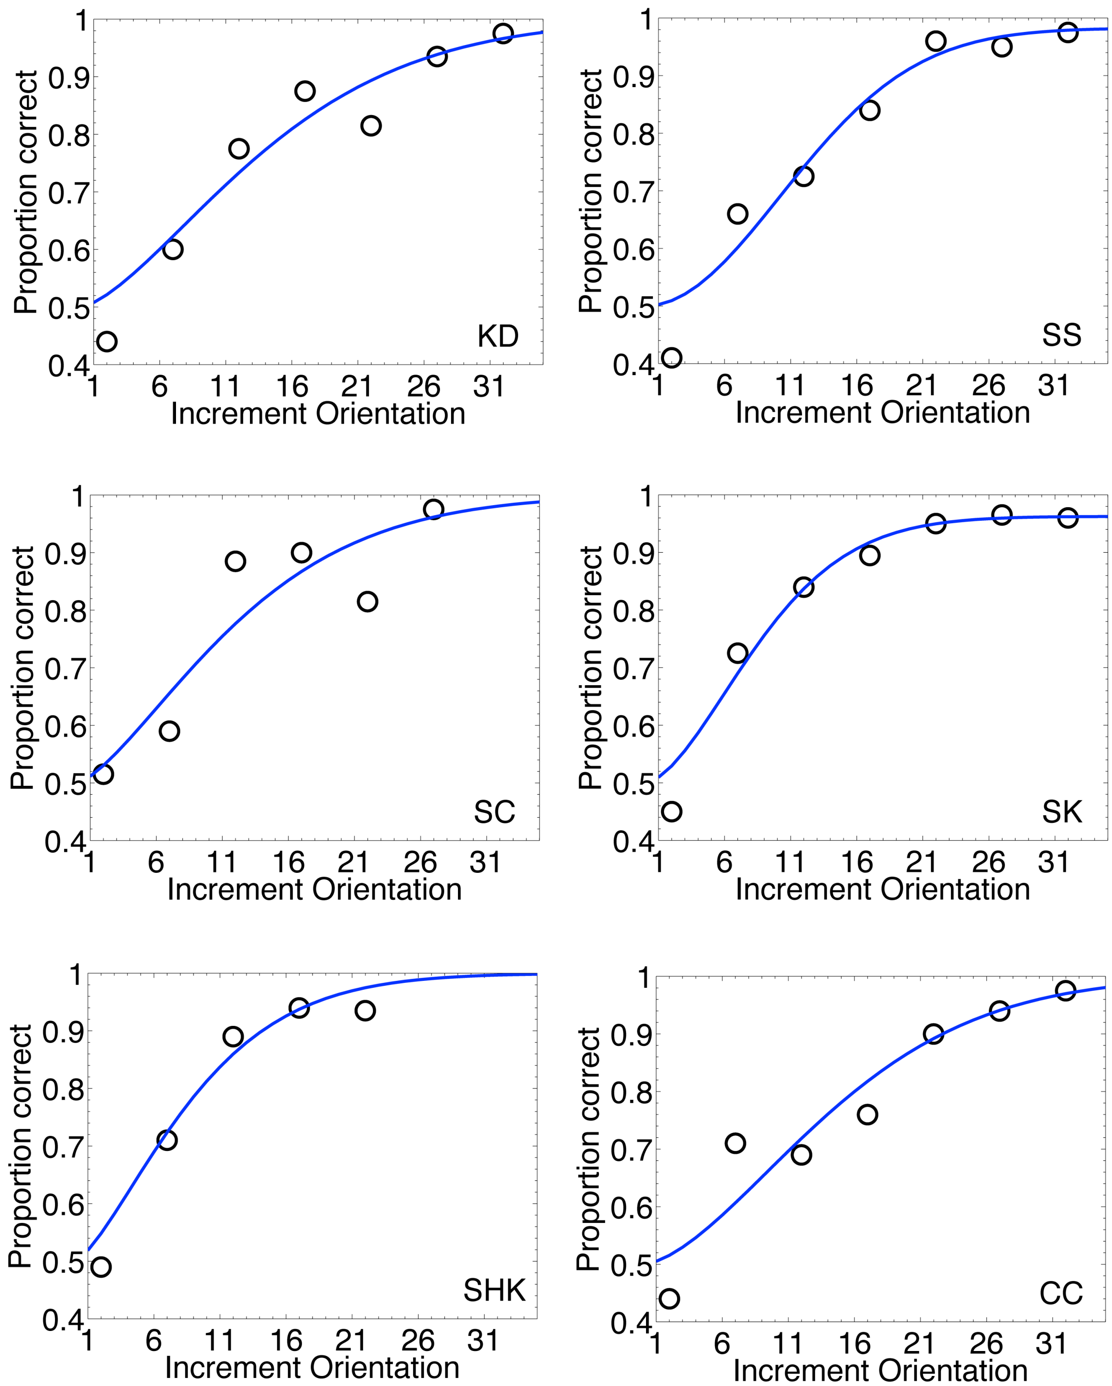

Figure 40b. Two-part shape (Experiment 3b): Random orientation transformation.

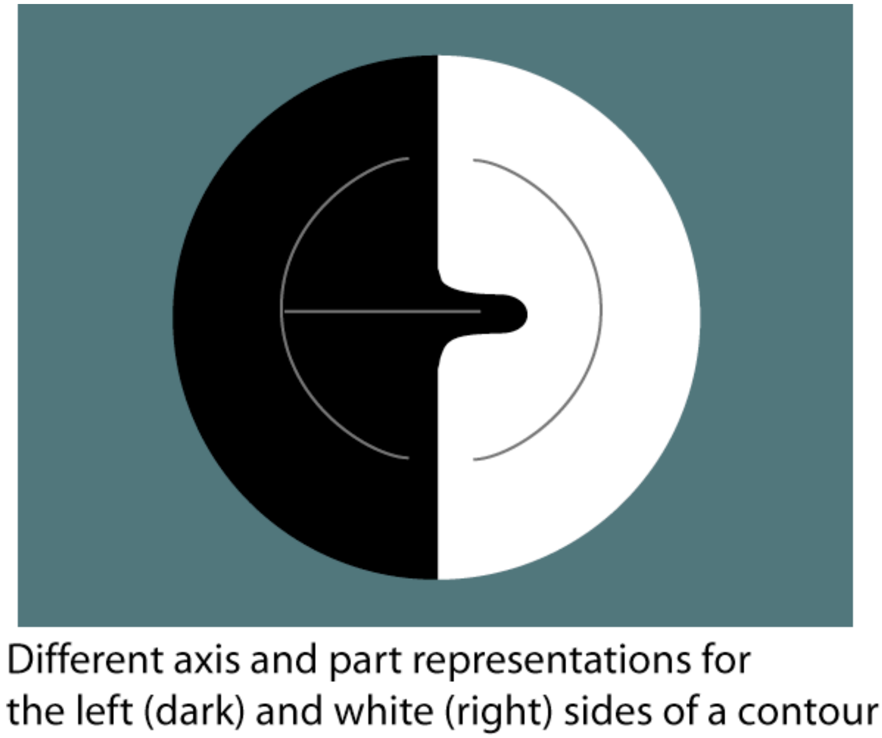

Figure 41. The figure-ground relationship. Given a single right-facing bulge contour, the surface on the right of the contour has an indentation and the surface on the left has a protrusion; both have different underlying axial structure.

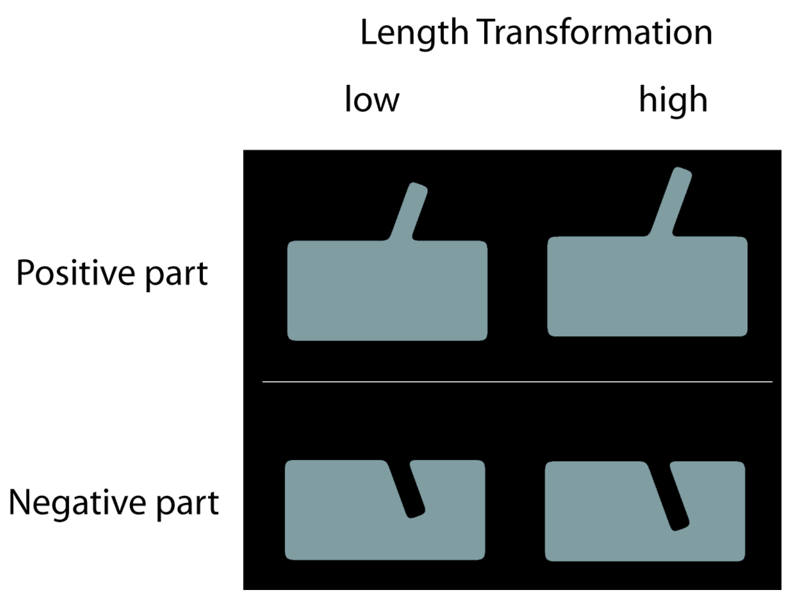

Figure 42. Two-part shape: wide base illustration. Top row illustrates low and high versions of transformations for the positive part; bottom row shows transformations for the negative part.

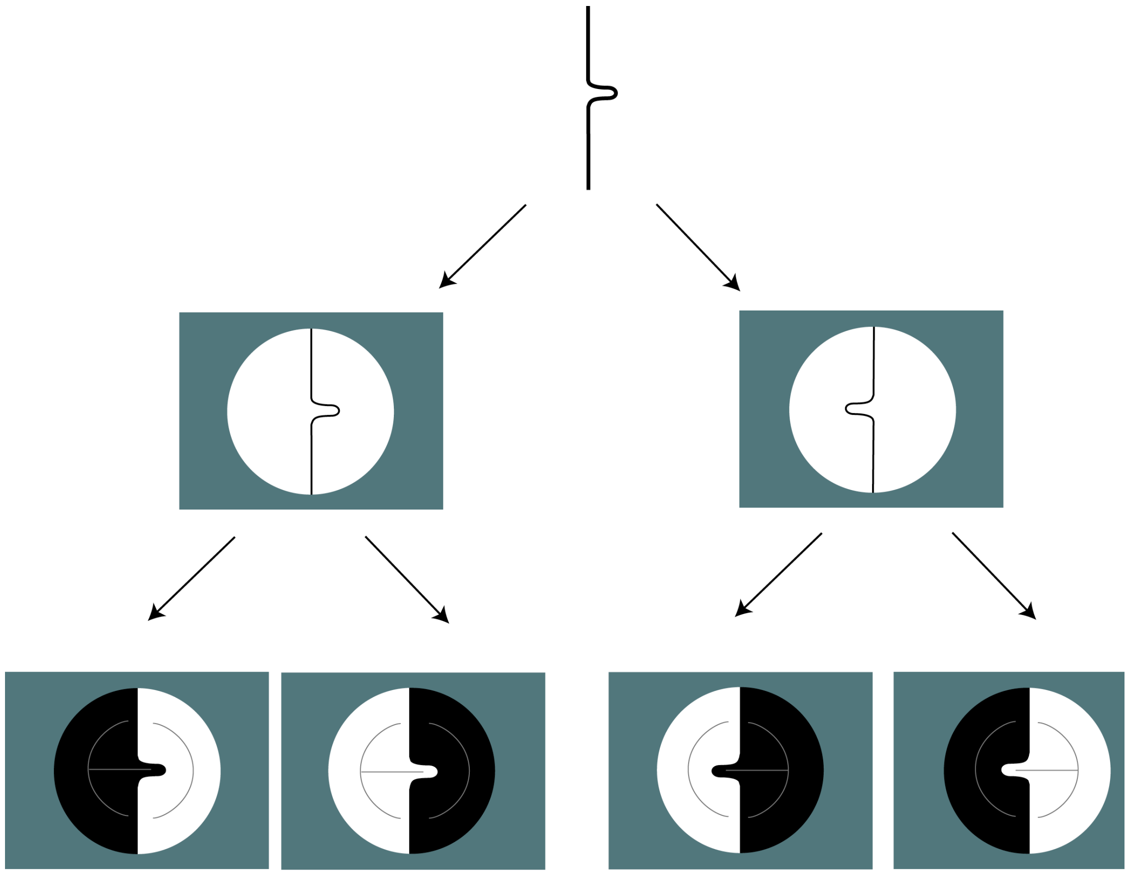

Figure 43. Illustration of contour- and surface- based shapes used in Experiment 4. A single contour, with its bulge facing either left or right, can give rise to two different types of surfaces depending on figure and ground relationships (See text for more details).

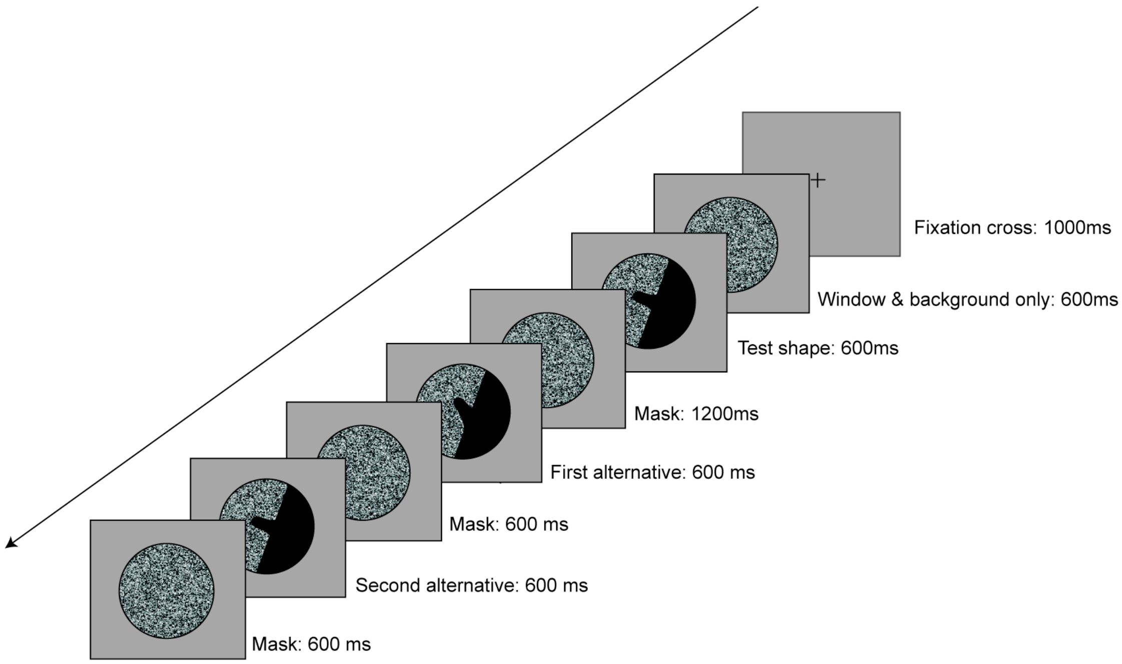

Figure 44. Trial sequence in Experiment 4. Note that all depth layers (the circular window, the shape, and the background) were presented as random dot textures. Different colors are used only for illustration purposes (i.e., the shape was not actually dark in the experimental sessions).

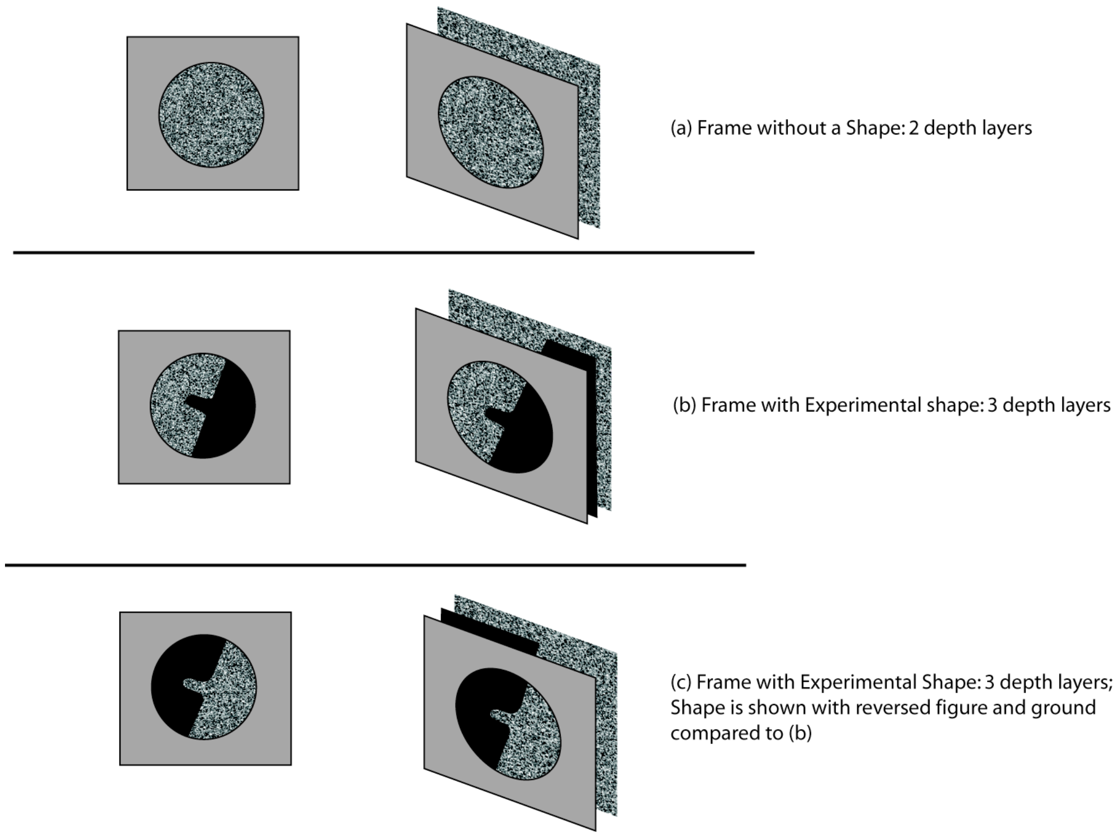

Figure 45. Illustration of depth relationships in different trial frames. Note that the depth layers in actual experimental sessions were presented as random dot textures (colors are used for illustrative purposes only).

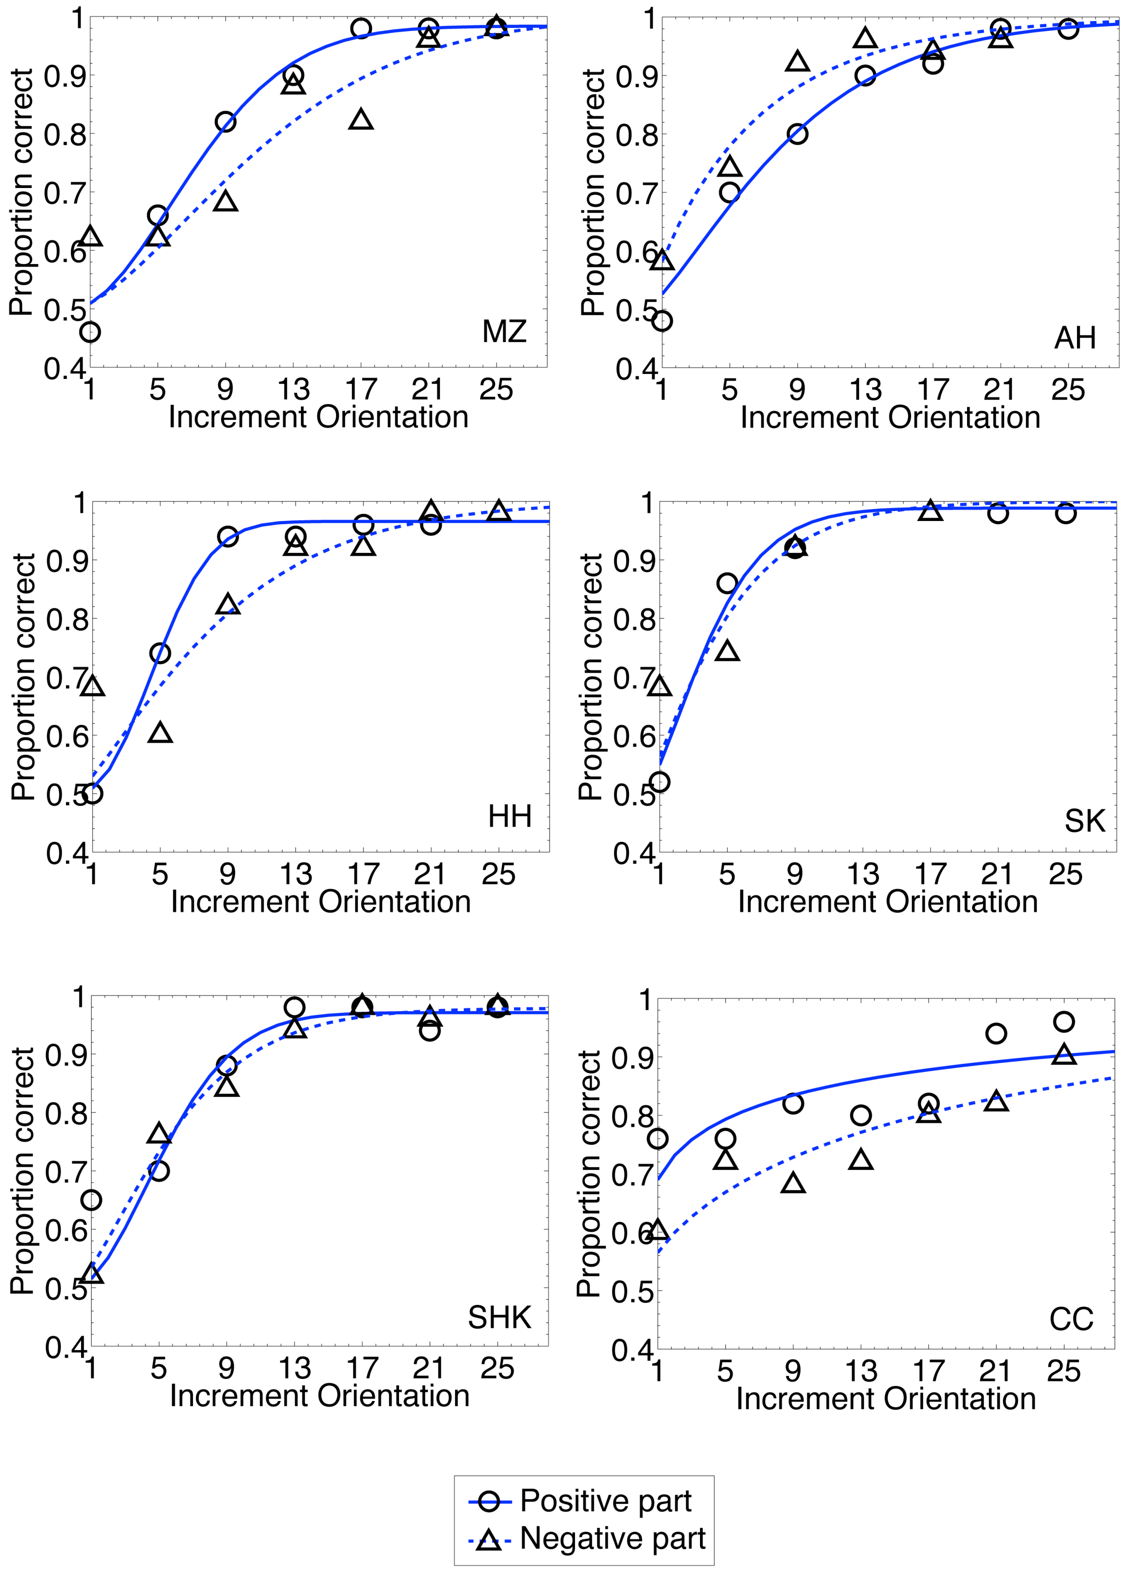

Figure 46a. Stereo presentation (Experiment 4): Orientation transformation. Observers are more sensitive to positive part compared to negative part transformations (Observer AH shows the opposite trend).

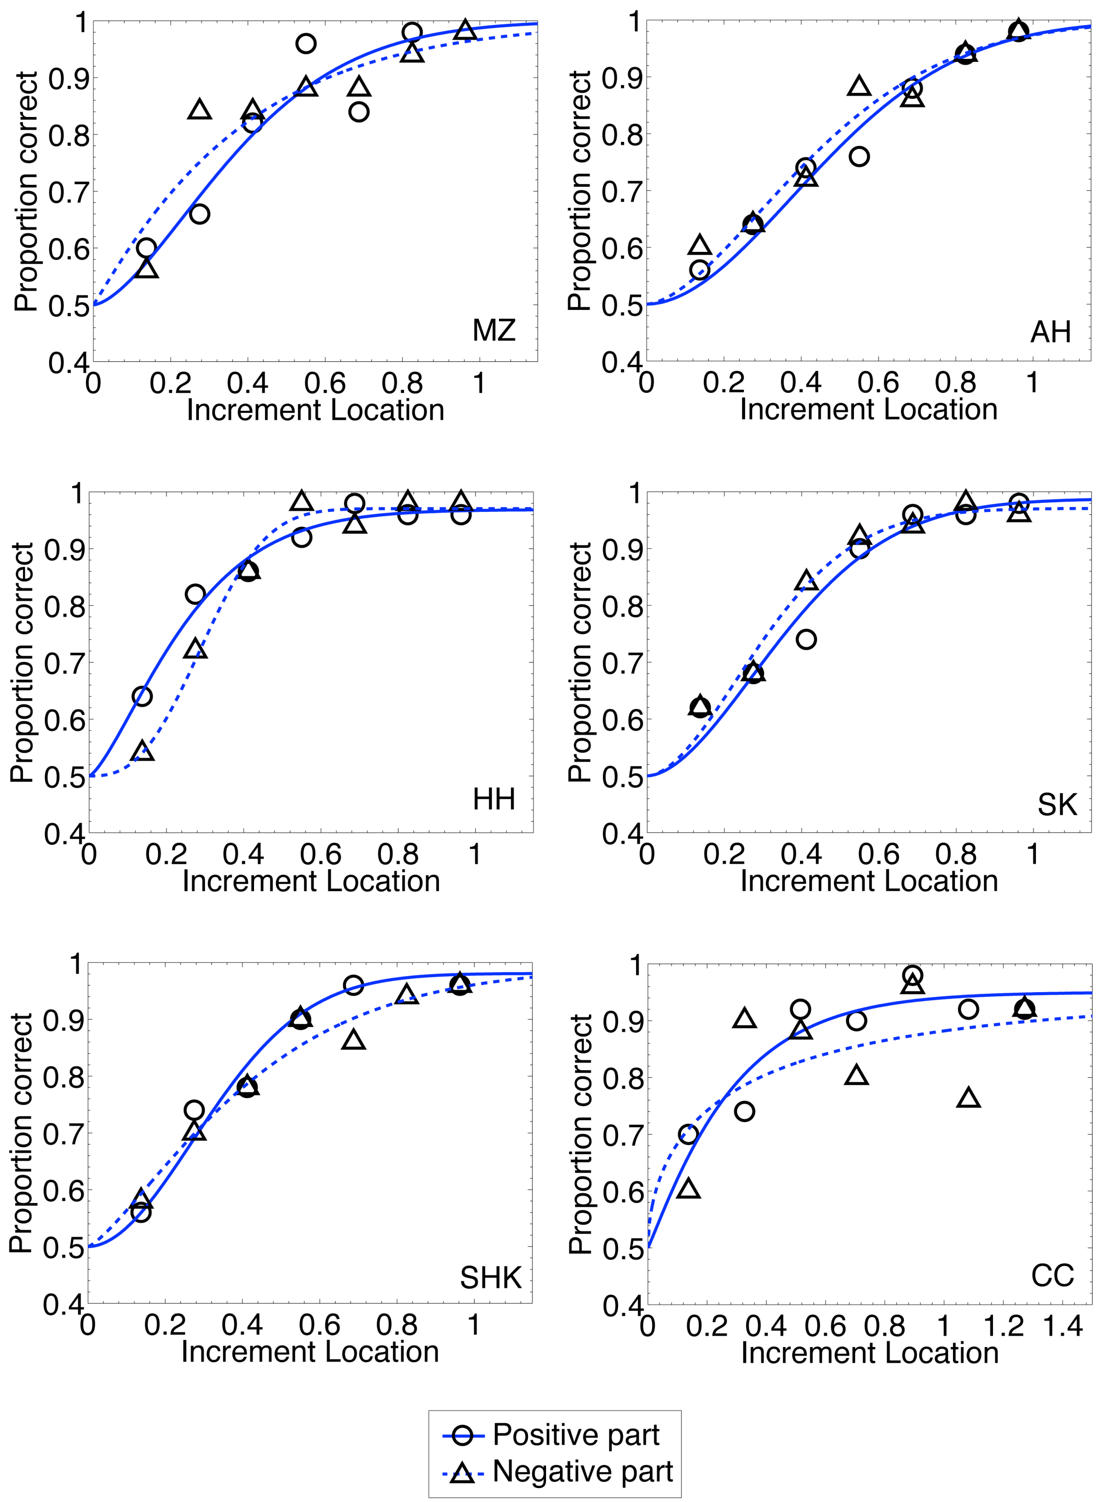

Figure 46b. Stereo presentation (Experiment 4b): Location transformation. Observers are consistently sensitive to neither to the positive nor to negative part transformations.

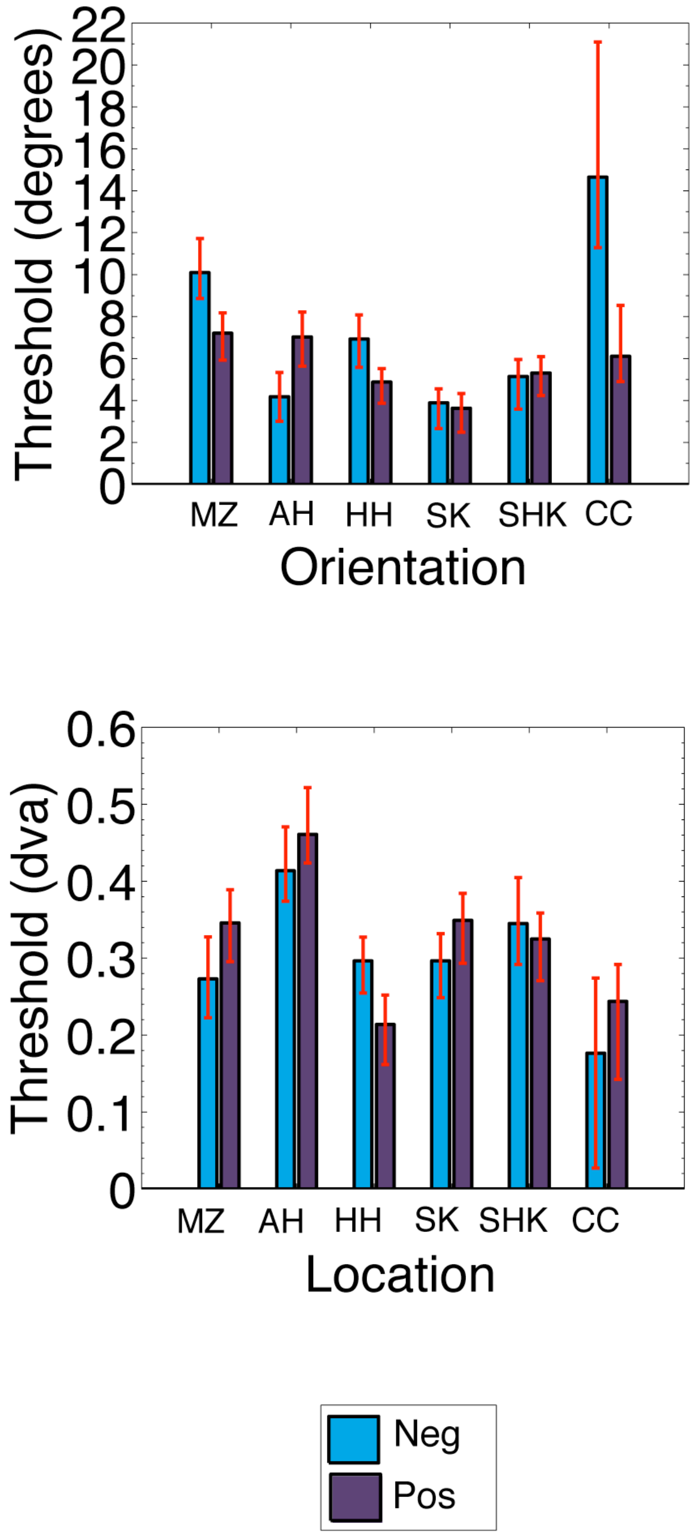

Figure 47. Experiment 4, Orientation & location (positive and negative transformations). Observers are more sensitive to positive compared to negative part changes in the orientation transformation; no pattern of sensitivity between positive and negative changes is observed in the location transformation.

Tables

	Raw threshold			Weber's fraction			Area difference			Distance (avg)		
KD	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.3032	0.2487	0.3297	0.1176	0.0965	0.1279	5.48	4.56	6.05	0.0392	0.0323	0.0421
width	0.1137	0.0987	0.1254	0.1322	0.1148	0.1459	6.19	5.51	6.82	0.0457	0.0399	0.0507
curv	0.1629	0.1454	0.1803	0.4000	0.3571	0.4429	19.84	17.83	21.67	0.1382	0.1238	0.1509
ori	10.7000	9.4660	11.9997	0.5350	0.4733	0.6000	27.32	24.15	30.55	0.1911	0.1690	0.2142
loc	0.5898	0.4639	0.6334	0.1072	0.0843	0.1151	65.29	51.67	69.96	0.3945	0.3709	0.4015
SS	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.2205	0.1580	0.2490	0.0855	0.0613	0.0966	4.20	3.10	4.56	0.0301	0.0206	0.0323
width	0.0711	0.0564	0.0790	0.0827	0.0656	0.0919	4.09	3.29	4.60	0.0290	0.0227	0.0329
curv	0.1222	0.1047	0.1338	0.3000	0.2571	0.3286	15.28	13.26	16.61	0.1060	0.0919	0.1149
ori	6.0800	4.8939	6.9570	0.3040	0.2447	0.3479	15.50	12.62	17.78	0.1088	0.0880	0.1244
loc	0.6190	0.5779	0.6881	0.1125	0.1051	0.1251	68.34	64.11	75.77	0.3979	0.3938	0.4088
SC	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.2140	0.1663	0.2435	0.0830	0.0645	0.0944	4.12	3.24	4.51	0.0271	0.0237	0.0323
width	0.0601	0.0484	0.0681	0.0699	0.0564	0.0792	3.44	2.86	3.81	0.0240	0.0200	0.0271
curv	0.0931	0.0814	0.1047	0.2286	0.2000	0.2571	11.83	10.50	13.26	0.0816	0.0722	0.0919
ori	9.4500	8.1884	10.4195	0.4725	0.4094	0.5210	24.11	20.87	26.58	0.1687	0.1463	0.1860
loc	0.3741	0.3097	0.4218	0.0680	0.0563	0.0767	41.95	34.78	47.05	0.3037	0.2501	0.3400
SK	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.3011	0.2774	0.3585	0.1168	0.1076	0.1390	5.44	5.16	6.39	0.0392	0.0362	0.0449
width	0.0593	0.0468	0.0672	0.0690	0.0545	0.0781	3.40	2.76	3.79	0.0237	0.0194	0.0268
curv	0.0989	0.0873	0.1163	0.2429	0.2143	0.2857	12.52	11.15	14.68	0.0864	0.0769	0.1014
ori	6.7616	5.6154	7.5901	0.3381	0.2808	0.3795	17.22	14.42	19.37	0.1209	0.1010	0.1356
loc	0.2353	0.1348	0.2730	0.0428	0.0245	0.0496	26.65	15.38	30.77	0.1915	0.1096	0.2216
SHK	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.2342	0.1935	0.2554	0.0908	0.0750	0.0990	4.39	3.62	4.67	0.0323	0.0260	0.0333
width	0.0526	0.0400	0.0600	0.0612	0.0466	0.0698	3.04	2.32	3.44	0.0214	0.0159	0.0240
curv	0.0989	0.0873	0.1105	0.2429	0.2143	0.2714	12.52	11.15	14.02	0.0864	0.0769	0.0969
ori	10.4321	9.2330	11.2178	0.5216	0.4617	0.5609	26.60	23.52	28.61	0.1862	0.1648	0.2005
loc	0.2416	0.1719	0.2751	0.0439	0.0313	0.0500	27.30	19.55	31.05	0.1962	0.1393	0.2232
CC	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.3922	0.3407	0.4198	0.1521	0.1321	0.1628	7.11	6.22	7.64	0.0496	0.0441	0.0504
width	0.0932	0.0808	0.1107	0.1084	0.0940	0.1288	5.23	4.66	6.05	0.0380	0.0336	0.0445
curv	0.1047	0.0814	0.1163	0.2571	0.2000	0.2857	13.26	10.50	14.68	0.0919	0.0722	0.1014
ori	10.1415	8.6745	11.1635	0.5071	0.4337	0.5582	25.84	22.13	28.47	0.1809	0.1549	0.1996
loc	0.3474	0.3021	0.4088	0.0631	0.0549	0.0743	39.01	34.00	45.66	0.2817	0.2446	0.3306

Table 1. Experiment 1: summary of raw thresholds, Weber's fractions, Area difference metric, and Average distance metric and corresponding CIs for length, width, orientation, and curvature for the two-part shape, for six observers. Length, width, and location thresholds are expressed in degree of visual angle (dva) and curvature thresholds in 1/dva. Orientation thresholds are shown in degrees. Webers fractions are computed by using the formula: $\Delta I / I$, where ΔI refers to the threshold value, and I is the baseline ("test") value for a given transformation. The baseline value for length was 2.58 dva; width: .86 dva; and curvature: .41 1/dva.

	Raw threshold			Weber's fraction			Area difference			Distance (avg)		
KD	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.3930	0.3690	0.4531	0.1524	0.1431	0.1757	7.78	7.57	8.91	0.0478	0.0464	0.0543
width	0.1088	0.0984	0.1203	0.1266	0.1145	0.1399	6.35	5.85	6.90	0.0430	0.0393	0.0472
curv	0.1454	0.1163	0.1571	0.3571	0.2857	0.3857	15.30	12.64	16.57	0.1220	0.0998	0.1314
ori	12.7982	9.9907	13.7344	0.6399	0.4995	0.6867	25.71	19.87	27.56	0.2283	0.1796	0.2420
SS	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.2122	0.1709	0.2341	0.0823	0.0663	0.0908	4.51	3.64	4.85	0.0276	0.0227	0.0307
width	0.0704	0.0597	0.0790	0.0819	0.0695	0.0919	4.54	3.62	4.95	0.0291	0.0239	0.0322
curv	0.0698	0.0565	0.0756	0.1714	0.1387	0.1857	7.77	6.18	8.34	0.0616	0.0497	0.0664
ori	5.9236	4.2825	6.8730	0.2962	0.2141	0.3437	11.97	8.60	13.76	0.1068	0.0775	0.1238
SC	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.1851	0.1512	0.2081	0.0718	0.0587	0.0807	3.87	3.04	4.49	0.0249	0.0203	0.0279
width	0.0833	0.0716	0.0927	0.0969	0.0833	0.1078	5.12	4.60	5.56	0.0337	0.0296	0.0372
curv	0.0756	0.0640	0.0873	0.1857	0.1571	0.2143	8.34	7.16	9.52	0.0664	0.0569	0.0758
ori	5.1465	3.9887	5.7585	0.2573	0.1994	0.2879	10.29	8.04	11.66	0.0929	0.0721	0.1039
SK	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.1413	0.1106	0.1608	0.0548	0.0429	0.0624	2.91	2.43	3.53	0.0194	0.0157	0.0214
width	0.0957	0.0850	0.1035	0.1114	0.0989	0.1204	5.70	5.21	6.12	0.0382	0.0344	0.0408
curv	0.0640	0.0529	0.0756	0.1571	0.1299	0.1857	7.16	5.78	8.34	0.0569	0.0466	0.0664
ori	5.7272	4.6141	6.8045	0.2864	0.2307	0.3402	11.59	9.26	13.63	0.1033	0.0834	0.1226
SHK	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.1583	0.1275	0.1806	0.0614	0.0495	0.0701	3.49	2.69	3.83	0.0219	0.0181	0.0238
width	0.0665	0.0530	0.0760	0.0774	0.0616	0.0885	4.02	3.33	4.79	0.0263	0.0212	0.0309
curv	0.0640	0.0483	0.0698	0.1571	0.1185	0.1714	7.16	5.37	7.77	0.0569	0.0425	0.0616
ori	7.1017	5.7632	8.5041	0.3551	0.2882	0.4252	14.20	11.66	17.00	0.1279	0.1040	0.1530
CC	T	CI (low)	CI (high)	W	CI (low)	CI (high)	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	0.2995	0.2697	0.3426	0.1162	0.1046	0.1329	6.03	5.59	6.96	0.0378	0.0348	0.0430
width	0.0869	0.0715	0.0998	0.1011	0.0832	0.1161	5.30	4.58	5.92	0.0351	0.0295	0.0398
curv	0.0756	0.0551	0.0931	0.1857	0.1354	0.2286	8.34	6.05	10.17	0.0664	0.0485	0.0804
ori	6.1545	4.6414	7.1730	0.3077	0.2321	0.3587	12.43	9.33	14.33	0.1110	0.0839	0.1292

Table 2. Experiment 2: summary of raw thresholds, Weber's fractions, Area difference metric, and Average distance metric and corresponding CIs for length, width, orientation, and curvature for the single-axis shape, for six observers. Length and width thresholds are expressed in degree of visual angle (dva) and curvature thresholds in 1/dva. Orientation thresholds are shown in degrees.

	Area difference			Distance (avg)		
KD	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	7.20	6.84	8.15	0.0486	0.0469	0.0561
width	6.26	5.67	6.78	0.0430	0.0393	0.0472
curv	14.38	11.90	15.72	0.0953	0.0780	0.1038
ori	15.73	12.37	16.81	0.1038	0.0817	0.1110
SS	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	4.11	3.30	4.42	0.0285	0.0220	0.0310
width	4.34	3.54	4.72	0.0291	0.0239	0.0322
curv	7.28	5.83	7.77	0.0475	0.0381	0.0513
ori	7.39	5.39	8.43	0.0491	0.0356	0.0565
SC	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	3.53	2.78	4.09	0.0235	0.0198	0.0274
width	4.95	4.37	5.36	0.0337	0.0296	0.0372
curv	7.77	6.70	8.94	0.0513	0.0437	0.0587
ori	6.40	4.99	7.23	0.0429	0.0335	0.0475
SK	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	2.62	2.23	3.13	0.0179	0.0154	0.0221
width	5.54	5.02	5.98	0.0382	0.0344	0.0408
curv	6.70	5.52	7.77	0.0437	0.0357	0.0513
ori	7.19	5.80	8.35	0.0472	0.0381	0.0560
SHK	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	3.09	2.45	3.41	0.0221	0.0175	0.0249
width	3.88	3.25	4.60	0.0263	0.0212	0.0309
curv	6.70	5.11	7.28	0.0437	0.0326	0.0475
ori	8.79	7.23	10.50	0.0586	0.0475	0.0699
CC	A	CI (low)	CI (high)	D	CI (low)	CI (high)
length	5.56	5.09	6.34	0.0390	0.0358	0.0431
width	5.08	4.37	5.77	0.0351	0.0295	0.0398
curv	7.77	5.71	9.42	0.0513	0.0372	0.0623
ori	7.62	5.82	8.89	0.0508	0.0384	0.0590

Table 3. Experiment 2 (maximally aligned shapes): summary of Area difference metric, and Average distance metric and corresponding CIs for length, width, orientation, and curvature for the single-axis shape, for six observers. Length and width thresholds are expressed in degree of visual angle (dva) and curvature thresholds in 1/dva. Orientation thresholds are shown in degrees.

(a)

	Raw threshold: orientation		
	T	CI (low)	CI (high)
KD	21.7755	20.5981	24.6144
SS	7.6415	6.2546	8.7691
SC	7.3171	5.852	8.8567
SK	8.1685	6.5868	9.4225
SHK	6.2283	4.6253	7.1892
CC	11.7784	10.2155	13.5718

(b)

	Raw threshold: orientation		
	T	CI (low)	CI (high)
KD	12.8151	11.3462	14.9305
SS	11.9924	10.5934	13.7008
SC	10.7993	9.2202	12.4836
SK	8.2233	6.4644	9.3598
SHK	7.8680	6.3217	9.0741
CC	13.4733	11.8193	15.5724

Table 4. Experiment 3: summary of thresholds and CIs for random orientation condition, for six observers. (a) shows thresholds for the single-axis shape and (b) shows thresholds for the two-part shape (b). Thresholds are shown in degrees.

Raw threshold

MZ	T	CI (low)	CI (high)
ori (N)	10.1005	8.8652	11.7178
ori (P)	7.2044	5.9273	8.1742
location(N)	0.2731	0.2226	0.3277
location(P)	0.3458	0.2955	0.3890

AK	T	CI (low)	CI (high)
ori (N)	4.1841	3.0128	5.3398
ori (P)	7.0300	5.6390	8.2133
location(N)	0.4136	0.3740	0.4707
location(P)	0.4609	0.4238	0.5218

HH	T	CI (low)	CI (high)
ori (N)	6.9313	5.5827	8.0779
ori (P)	4.8807	3.8695	5.5198
location(N)	0.2964	0.2546	0.3273
location(P)	0.2137	0.1617	0.2520

SK	T	CI (low)	CI (high)
ori (N)	3.8886	2.6600	4.5601
ori (P)	3.6330	2.4923	4.3307
location(N)	0.2967	0.2487	0.3318
location(P)	0.3492	0.2935	0.3842

SHK	T	CI (low)	CI (high)
ori (N)	5.1453	3.5916	5.9541
ori (P)	5.3072	4.2386	6.0962
location(N)	0.3451	0.2918	0.4049
location(P)	0.3249	0.2708	0.3589

CC	T	CI (low)	CI (high)
ori (N)	14.6469	11.2913	21.0889
ori (P)	6.1055	4.8988	8.5340
location(N)	0.1765	0.0270	0.2742
location(P)	0.2437	0.1425	0.2917

Table 5. Experiment 4: summary of thresholds and CIs for orientation and location transformations, for six observers. Location thresholds are shown in dva; orientation thresholds are shown in degrees. Positive transformations are designated (P), and negative transformations are designated (N).

References

Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, 61, 183-193.

August, J., Siddiqi, K., & Zucker, S.W. (1999). Ligature instabilities in the perceptual organization of shape. *Computer Vision and Image Understanding*, 76(3), 231-243.

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. (1995). One-sided edge assignment in vision. 1. Figure-ground segmentation and attention to objects. *Current Directions in Psychological Science*, 4, 140-146.

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

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

Barenholtz, E., & Feldman, J. (2006). Determination of visual figure and ground in dynamically deforming shapes. *Cognition*, 101, 530-544.

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

Bertamini, M., & Farrant, T. (2005). Detection of change in shape and its relation to part structure, *Acta Psychologica*, 120, 35-54.

Bertamini, M., & Lawson, R. (2008). Rapid figure-ground responses to stereograms reveal an advantage for a convex foreground. *Perception*, 37, 483-494.

Biederman, I. (1985). Human image understanding: Recent research and a theory. *Computer Vision, Graphics, and Image Understanding*, 32, 29-73.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94(2), 115-117.

Binford, T.O. (1971). *Visual perception by computer*. Paper presented at the IEEE Conference on Systems and Control, Miami, FL.

Blum, H. (1973). Biological shape and visual science (Part I). *Journal of Theoretical Biology*, 38, 205-287.

Briscoe, E. J. (2009). Shape skeletons and shape similarity. *Dissertation Abstracts International: Section B: The Sciences and Engineering*. 70(3-B), 1962.

Burbeck, C. A., & Regan, D. (1983). Independence of orientation and size in spatial discriminations, *J. Opt. Soc. Am.*, 73, 1691-1694.

Burbeck, C. A., & Pizer, S. M. (1995). Object representation by cores: Identifying and representing primitive spatial regions. *Vision Research*, 35, 1917-1030.

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.

Cohen, E. H., & Singh, M. (2006). Perceived orientation of complex shape reflects

graded part decomposition. *Journal of Vision*, 6, 805-821.
Cohen, E. H., & Singh, M. (2007). Geometric determinants of shape segmentation: Tests using segment identification. *Vision Research*, 47, 2825-2840.
Denisova, K., Singh, M., & Kowler, E. (2006). The role of part structure in the perceptual localization of a shape. *Perception*, 35, 1073-1087.
De Winter, J., & Wagemans, J. (2006). Segmentation of object outlines into parts: A large-scale integrative study. *Cognition*, 99, 275-325.
Droll, A.D., & Hayhoe, M. (2007). Trade-offs between gaze and working memory use. *Journal of Experimental Psychology: Human Perception and Performance*, 33(6), 1352–1365.
Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, 113(4), 501-517.
Feldman, J., & Singh, M. (2005). Information along contours and object boundaries. *Psychological Review*, 112, 243-252.
Feldman, J., & Singh, M. (2006). Bayesian estimation of the shape skeleton. *Proceedings of the National Academy of Sciences, USA*, 103(47), 18014-18019.
Ginsburg, A. P. (1971). Psychological correlates of a model of the human visual system. Paper presented at the *IEEE Transactions on Aerospace and Electronic Systems*, New York.
Ginsburg, A. P. (1986). Spatial filtering and visual form perception. In K.R.Boff, L. Kaufman, & J.P. Thomas (Eds.) *Handbook of perception and human performance: Vol. 2. Cognitive processes and performance* (pp.1-41). New York: Wiley.
Gray, R., & Regan, D. (1996). Accuracy of reproducing angles: is a right angle special? *Perception*, 25(5), 531-542.
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.
Hulleman, J., te Winkel, W., & Boselie, F. (2000). Concavities as basic features in visual search: Evidence from search asymmetries. *Perception & Psychophysics*, 62, 162-174.
Julesz, B. (1971). *Foundations of cyclopean perception*. Chicago: U of Chicago Press.
Kabrisky, M. (1966). *A proposed model for visual information processing in the human brain*. Urbana, IL: University of Illinois Press.
Kanizsa, G., & Gerbino, W. (1976). Convexity and symmetry in figure-ground organization, in *Art and Artefacts*. M. Henle (Ed.) New York: Springer, pp. 25-32.
Keane, S.K., Hayward, W.G., & Burke, D. (2003). Detection of three types of changes to novel objects. *Visual Cognition*, 10 (1), 101-127.
Kennedy, J. M., & Domander, R. (1985). Shape and contour: The points of maximum change are least useful for recognition. *Perception*, 14, 367-370.
Kimchi, R., & Bloch, B. (1998). Dominance of configural properties in visual form perception. *Psychonomic Bulletin and Review*, 5 (1), 135-139.
Kovacs, I., & Julesz, B. (1994). Perceptual sensitivity maps within globally defined visual shapes. *Nature*, 370, 644-646.
Kovacs, I., Feher, A., & Julesz, B. (1998). Medial-point description of shape: A representation for action coding and its psychophysical correlates. *Vision Research*, 38, 2323-2333.

Latecki, L., & Lakamper, R. (2000). Shape similarity measure based on correspondence of visual parts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(10), 1-6.

Leyton, M. (1992). *Symmetry, causality, mind*. Cambridge, MA: MIT Press.

Leyton, M. (1988). A process grammar for shape. *Cognitive Science*, 13, 213-247.

Leyton, M. (1989). Inferring causal history from shape. *Cognitive Science*, 13, 357-387.

Lipton, L. (2001). The stereoscopic cinema: from film to digital projection. *SMPTE Journal*, 586-593.

Mach, E. (1914/1959). *The analysis of sensations*. Chicago: Open Court.

Marr, D. & Nishihara, H.K. (1978). Representation and recognition of the spatial organization of three-dimensional shapes. *Proceedings of the Royal Society of London*, 200, 269-294.

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W. H. Freeman.

Melcher, D. (2006). Accumulation and persistence of memory for natural scenes. *Journal of Vision*, 6, 8–17.

Mi, X., DeCarlo, D., & Stone, M. (2009). Abstraction of 2D shapes in terms of parts. *NPAR*, 15-24.

Mi, X., & DeCarlo, D. (2007). Separating parts from 2D shapes using relatability. *ICCV*, 1-8.

Norman, J.F., Phillips, F., & Ross, H. E. (2001). Information concentration along the boundary contours of naturally shaped solid objects. *Perception*, 30, 1285-1294.

Palmer, S. E. (2001). *Vision Science: Photons to Phenomenology*. Cambridge, Massachusetts: MIT Press.

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

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

Pasupathy, A., & Connor, C. E. (2002). Population coding of shape in are V4. *Nature Neuroscience*, 5(12), 1332-1338.

Psotka, J. (1978). Perceptual processes that may create stick figures and balance. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 101-111.

Regan, D., & Hamstra, S. J. (1992). Shape discrimination and the judgment of perfect symmetry: dissociation of shape from size. *Vision Research*, 32, 1845-1864.

Rensink, R. A., O'Regan, K. A., & Clark, J. J. (1997). The need for attention to perceive changes to scenes. *Psychological Science*, 8, 368-373.

Rosin, P. (2000). Shape partitioning by convexity. *IEEE Transactions on Systems, Man, and Cybernetics—Part A. Systems and humans,* 30(2), 202-210.

Sabra, A. I. (1989). *The optics of Ibn al-Haytham*. London: Warburg Institute.

Sebastian, T.B., & Kimia, B.B. (2005). Curves vs. skeletons in object recognition. *Signal Processing*, 85, 247-263.

Simons, D. J., & Levin, D. (1997). Change blindness. *Trends in Cognitive Sciences*, 1, 261-267.

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 in vision, advances in psychology* (Vol. 130, pp.401-459). New York: Elsevier.

Shepard, R. N. (1962a). The analysis of proximities: Multidimensional scaling with an unknown distance function: Part I. *Psychometrika*, 27, 125-140.

Shepard, R. N. (1962b). The analysis of proximities: Multidimensional scaling with an unknown distance function: Part II. *Psychometrika*, 27(3), 219-246.

Stevens, K. A., & Brookes, A., (1988). The concave cusp as a determiner of figure-ground. *Perception,* 17(1) 35-42.

Supèr, H., Spekreijse, H., & Lamme, V.A.F. (2001). A neural correlate of working memory in the monkey primary visual cortex. *Science*, 293, 120-124.

Ullman, S. (1989). Aligning pictorial descriptions: An approach to object recognition. *Cognition*, 32, 193-254.

Watt R. J., & Andrews, D. P. (1982). Contour curvature analysis: hyperacuities in the discrimination of detailed shape. *Vision Research*, 22, 449-460.

Watson, S.E., & Kramer, A.F. (1999). Object-based visual selective attention and perceptual organization. *Perception and Psychophysics*, 61(1), 31-49.

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

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

Zhang, L., & Cottrell, G. W. (2005). Holistic processing develops because it is good. In B.G. Bara, L. Barsalou, & M. Bucciarelli (Eds.) *Proceedings of the 27th annual cognitive science conference* (pp. 2428-2433). Mahwah, NJ: Erlbaum.

KRISTINA DENISOVA

Education

Ph.D., Cognitive Psychology, Rutgers University (October, 2010)

M.S., Cognitive Psychology, Rutgers University, New Brunswick, NJ (May, 2005)

Certificate in Cognitive Science, RUCCS: Rutgers University Center for Cognitive Science (October, 2010)

B.A., Psychology (magna cum laude), New York University, NY (May, 2001)

Academic Occupations

• NSF IGERT Fellow (*Integrative Graduate Education and Research Traineeship*) in Perceptual Science (Rutgers University, 2007-2009)

• Judge (Behavioral Science Category), North Jersey Regional Science Fair (Rutgers University, March 19, 2010)

• Judge (Behavioral Science Category), North Jersey Regional Science Fair (Rutgers University, March 20, 2009)

Publications

Carrasco, M., McElree, B., Denisova, K., & Giordano, A. (2003). "Speed of visual processing increases with eccentricity". *Nature Neuroscience,* **6**, 699 – 700.

Denisova K., Singh M., & Kowler, E. (2006) "The role of part structure in the perceptual localization of a shape". *Perception*, **35**(8), 1073 – 1087.