Available online at www.sciencedirect.com
ScienceDirect
Acta Psychologica 123 (2006) 20–40
acta psychologica
www.elsevier.com/locate/actpsy

Surface geometry influences the shape of illusory contours

Jacqueline M. Fulvio [a], Manish Singh [b,*]

[a] *Department of Psychology, New York University, USA*
[b] *Department of Psychology and Center for Cognitive Science, Rutgers, New Brunswick, USA*

Received 26 October 2005; received in revised form 19 February 2006; accepted 24 February 2006
Available online 8 May 2006

Abstract

Geometric and neural models of illusory-contour (IC) synthesis currently use only local contour geometry to derive the shape of ICs. Work on the visual representation of shape, by contrast, points to the importance of both contour and surface geometry. We investigated the influence of surface-based geometric factors on IC shape. The local geometry of inducing-contour pairs was equated in stereoscopic IC displays, and the shape of the enclosed surface was varied by manipulating sign of curvature, cross-axial shape width, and medial-axis geometry. IC shapes were measured using a parametric shape-adjustment task (Experiment 1) and a dot-adjustment task (Experiment 2). Both methods revealed large influences of surface geometry. ICs enclosing locally concave regions were perceived to be systematically more angular than those enclosing locally convex regions. Importantly, the influence of sign of curvature was modulated significantly by shape width and medial-axis geometry: IC shape difference between convex and concave inducers was greater for narrow shapes than wider ones, and greater for shapes with straight axis and symmetric contours (diamond versus bowtie), than those with curved axis and parallel contours (bent tubes). Even at the level of illusory "contours," there is a contribution of region-based geometry which is sensitive to nonlocal shape properties involving medial geometry and part decomposition. Models of IC synthesis must incorporate the role of nonlocal region-based geometric factors in a way that parallels their role in organizing visual shape representation more generally.

* Corresponding author. Tel.: +1 732 445 6714; fax: +1 732 445 6715.
E-mail address: manish@ruccs.rutgers.edu (M. Singh).

0001-6918/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.actpsy.2006.02.004

PsycINFO classification: 2323

Keywords: Contour interpolation; Visual completion; Illusory contours; Shape perception; Parts; Part-based representation; Curvature; Convexity; Medial axis

1. Introduction

The ubiquitous occurrence of occlusion in natural scenes poses a difficult challenge to the visual computation of objects. Objects can either be partly occluded by an interposed object, or be camouflaged by an underlying surface that happens to project the same color (a problem that is exacerbated in conditions of low illumination). Objects are also invariably self-occluded: a large portion of each opaque object—its "back" relative to the viewer—has no counterpart in the retinal images. Moreover, occlusion and disocclusion can take place dynamically, either due to the motion of one of the objects, or that of the observer's vantage point. Michotte's work stands out as the first unified treatment of the full range of completion phenomena in both static and kinematic displays—where the visual system fills in missing structure—and a systematic analysis of the image cues necessary to initiate processes of visual completion (Michotte, 1950; Michotte, Thinès, & Crabbé, 1991). To this day, Michotte's work continues to motivate research into these problems; the articles in this special issue attest to the long-standing influence of his insights (see Bertamini & Hulleman, 2006; Van Lier, de Wit, & Koning, 2006).

The current paper focuses on the context of "illusory" or "subjective" contours—vivid contours that can be perceived in image regions containing no local contrast (see, e.g., Fig. 1). Following Michotte et al., this form of interpolation is referred to as *modal* completion, to emphasize the fact that the percept has the same 'mode' as if image contrast were actually present: "... these additions present the same visual qualities (luminance and color) as the rest of the configuration ..." (Michotte et al., 1991). It is contrasted with *amodal* completion—the visual completion of partly occluded contours and surfaces—where no contrast is perceived along the visually interpolated contour, despite a compelling

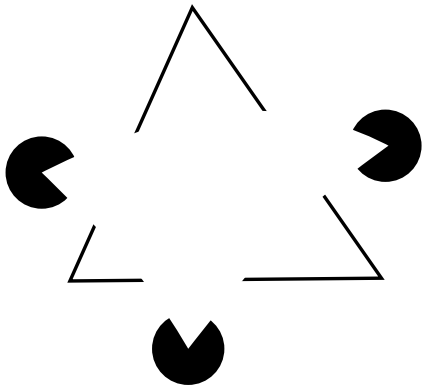

Fig. 1. An example of illusory contours. Vivid contours are perceived in image regions containing no local contrast.

percept of object unity. Although both forms generate the perception of completed contours, the perceptual vividness of illusory contours makes them particularly suitable for investigating visual mechanisms of shape completion.

Current geometric and neural models of illusory-contour synthesis take into account only the local geometry of inducing contours. Specifically, they encode interpolation geometry in terms of the relative positions of the physically specified "inducing" contours, and their relative orientations. The commonly used criterion of inducer *relatability*, for instance, is defined in terms of whether or not the linear extensions of the inducing contours intersect, and whether their interior angle of intersection is obtuse (Kellman & Shipley, 1991). (This criterion is equivalent to the existence of a smooth interpolating contour that contains no inflections, and undergoes a total turning of no more than 90°; Singh & Hoffman, 1999.) Based on these geometric relations, the relatability model predicts whether or not visual interpolation will occur. Similar geometric relations between pairs of oriented edges are also used to predict grouping in the context of contour integration (e.g., Elder & Goldberg, 2002; Field, Hayes, & Hess, 1993; Geisler, Perry, Super, & Gallogly, 2001; Neumann & Mingolla, 2001).

More directly relevant to the present study, models that predict specific shapes for illusory contours also take into account only contour geometry. These models vary a great deal in how they generate the interpolated shapes. An important class of models employs the minimization of a curvature-based functional—such as *total curvature* (Horn, 1983; Mumford, 1994) or *variation in curvature* (Kimia, Frankel, & Popescu, 2003). Of all interpolating shapes possible, such models take the preferred solution to be the one that minimizes either the squared curvature, or the squared derivative of curvature, integrated along the length of the interpolating contour. (See Singh & Fulvio, 2005 for an approach that models the shapes of visually extrapolated contours via a Bayesian interaction between the constraints of minimizing curvature and minimizing variation in curvature.) Other models build up the interpolated shape via propagation of local constraints, or local interactions between oriented edge pairs (Fantoni & Gerbino, 2003; Heitger, von der Heydt, Peterhans, Rosenthaler, & Kübler, 1998; Williams & Jacobs, 1997). Yet others postulate specific geometric forms for the interpolated contours, such as a pair of circular arcs (Ullman, 1976), or the combination of a circular arc and a line segment (Kellman & Shipley, 1991), that meet with continuous tangents.

Despite the different approaches adopted by these models, they share one important characteristic: they all formulate the problem of shape completion as one of *contour* interpolation. Specifically, interpolation geometry is encoded in terms of the local geometry of inducing-contour pairs, and the synthesis of the illusory shapes is driven by contour-based mechanisms—i.e., mechanisms involving 1D constraints applied along contours. Constraints based on the geometry of the enclosed region or surface are not considered. These models thus implicitly assume that region-based geometry plays no role in illusory-contour synthesis. In this article, we test this prevalent assumption: Are the shapes of illusory contours in fact determined simply by local contour geometry?

In the context of amodal completion, it has been shown that the perceived shape of a partly occluded object is often governed by global symmetry (e.g., radial or bilateral symmetry)—which results in a shorter description length for the occluded object—even if it entails a locally more complex shape for the amodally completed contour (De Wit & Van Lier, 2002; Sekuler, Palmer, & Flynn, 1994; Van Lier, Van der Helm, & Leeuwenberg, 1994, 1995; Van Lier, 1999). These cases of shape completion are theoretically important,

because they show that local contour geometry is not sufficient to predict the perceived shapes of partly occluded objects. However, such instances appear not to occur with illusory contours (see, e.g., Kellman, Guttman, & Wickens, 2001), which are the focus of the current study. More importantly, although this work makes a clear and important distinction between local and global geometry, it does not explicitly distinguish between contour and surface geometry. In its current form, for instance, the local computation of 'good continuation' and the computation of description length for the global shape are both contour based. (Shape is coded as a chain of contour elements and turning angles; although this scheme may of course be extended to incorporate surface geometry as well.)

Computational and psychophysical work on the visual representation of shape makes a clear distinction between a contour-based description and surface-based (or region-based) description of shape, and the contributions of both have been demonstrated in shape perception (see, e.g., Barenholtz & Feldman, 2003; Burbeck & Pizer, 1995; Sebastian & Kimia, 2005; Siddiqi, Tresness, & Kimia, 1996; Siddiqi, Kimia, Tannenbaum, & Zucker, 2001; Singh, Seyranian, & Hoffman, 1999). In the next section, we briefly review the role of three such region-based geometric factors in shape perception which will be relevant to our experiments: sign of curvature, medial-axis geometry, and cross-axial shape width.

2. Contour versus surface geometry in shape representation

In representing the shape of a closed contour, one may treat it either as a 1-D curve, or as a 2-D region (enclosed by the contour)—i.e., a "string" representation or "cardboard" representation. Importantly, the same local contour segment—e.g., the one shown in Fig. 2(a)—can bound surfaces with different region-based geometries. Four such examples are shown in Fig. 2(b)–(e). In Fig. 2(b), the contour segment corresponds to the convex protuberance of a single-part shape, in Fig. 2(c) it is a concavity corresponding to the boundary between two distinct parts, and in Fig. 2(d) and (e), it corresponds to the convex and concave side, respectively, of a bent tube.

2.1. Sign of curvature

Perhaps the most basic difference in region-based geometry arises from the border ownership of a contour, i.e., the side on which the figural surface lies that 'owns' the contour (Baylis & Driver, 1995; Koffka, 1935; Nakayama, He, & Shimojo, 1995).

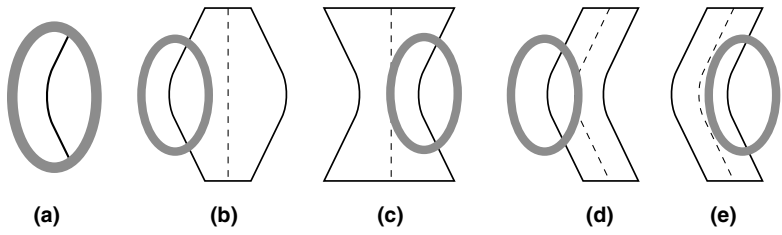

Fig. 2. Demonstrating the distinction between local contour geometry and nonlocal surface geometry. (a) A local contour segment viewed through an aperture. This segment can result from different surface geometries: (b) a protruberance on a convex object; (c) a part boundary between the two parts on a bowtie shape; (d) the convex side of a bent tube; and (e) the concave side of a bent tube.

Depending on border ownership, the same local contour segment may correspond either to a local concavity (hence with negative contour curvature) or to convexity (with positive curvature). The suggestion that local regions of positive and negative curvature play different roles in visual shape representation was made by Koenderink and van Doorn (1982), who noted that regions of positive curvature tend to have a thing-like character whereas regions of negative curvature tend to have a glue-like character. Since then, systematic differences based on the sign of curvature have been observed in various behavioral tasks including detecting small gaps in contours (Lamote & Wagemans, 1999), visual search (Hulleman, te Winkel, & Boselie, 2000), estimating vertex height (Bertamini, 2001), and change detection on shapes (Barenholtz, Cohen, Feldman, & Singh, 2003; Cohen, Barenholtz, Singh, & Feldman, 2005). An asymmetry in information content arising from the sign of contour curvature has been demonstrated using Shannon's definition of information (Feldman & Singh, 2005). Moreover, single-cell recordings in area V4 have shown that neurons in this area respond either to convex, or to concave, extrema of curvature—hence displaying a selectivity for the sign of curvature (Pasupathy & Connor, 1999).

Further evidence for a differential treatment of positive and negative curvature comes from the extensive work on part-based representation of shape (Barenholtz & Feldman, 2003; Bertamini & Farrant, 2005; Braunstein, Hoffman, & Saidpour, 1989; Baylis & Driver, 1994, 1995; De Winter & Wagemans, 2004, 2006; Hoffman & Richards, 1984; Hoffman & Singh, 1997; Lamote & Wagemans, 1999; Singh et al., 1999; Xu & Singh, 2002). Consistent with Hoffman and Richards' *minima rule*, negative minima of curvature (extrema of curvature in concave regions) are used by the visual system to define candidate boundaries between distinct parts, whereas corresponding extrema in convex regions typically do not play such a role (see Singh & Hoffman, 2001 for a review). An example is shown in Fig. 3, where the same undulating contour appears differently shaped on the two half disks (Attneave, 1971). This perceptual difference is readily explained by the switch in the sign of curvature that accompanies a figure-ground reversal—which in turn alters the locations of the negative minima, and hence part decomposition of the contour.

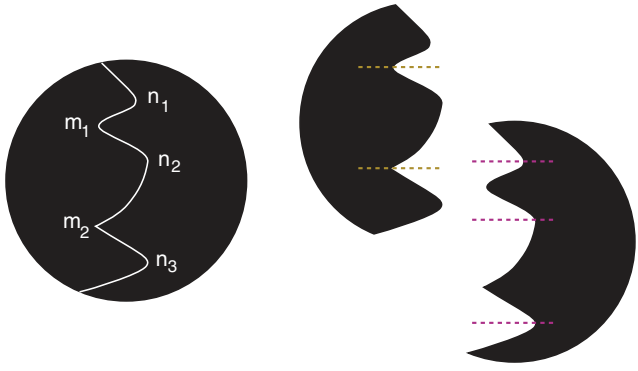

Fig. 3. Demonstrating the influence of sign of curvature on shape perception. The same undulating contour appears differently shaped on the two half disks. Changing the figural sidedness of a contour switches its sign of curvature. As a result, the negative minima of curvature that define part boundaries switch from locations m_1 and m_2 on the left half-disk to the locations n_1, n_2, and n_3 on the right half-disk.

2.2. Medial axis

Although the sign of curvature is determined by an aspect of surface geometry (namely, the surface *sidedness* of a contour), it is nevertheless entirely a local property. Indeed, the same local convexity or concavity may play an altogether different role in a shape's representation, depending on more global aspects of its geometry. An important such factor is a shape's "skeletal-axis" structure, as captured by the medial-axis representation. The geometric construction of the medial axis involves expressing a shape as a union of maximal disks inscribed within the shape (Blum, 1973). The shape is then represented in terms of (i) the medial axis, defined as the locus of the centers of the maximal disks, and (ii) the thickness function, defined by the (generally varying) diameters of the maximal disks along the medial axis. Variants of this construction have also been proposed (Brady & Asada, 1984; Leyton, 1988), including more recent algorithms based on curve evolution (Kimia, Tannenbaum, & Zucker, 1995; Siddiqi, Kimia, Tannenbaum, & Zucker, 1999).

Importantly, the same curved boundary segment can result from very different medial-axis geometries. In Fig. 4(a), the curvature of the bounding contour arises from systematic variations in the thickness of the shape about a straight medial axis, whereas in Fig. 4(b) its curvature arises entirely from the curvature of the axis itself (the thickness of the shape being constant). As a result, the shape in Fig. 4(a) is perceived as consisting of a number of distinct parts separated by narrow 'necks,' whereas the shape in Fig. 4(b) is perceived as a bent tube, or a 'snake,' with no distinct structural parts (Siddiqi et al., 1999, 2001). This difference has important psychophysical implications. For the 'peanut' shape in Fig. 4(a), Barenholtz and Feldman (2003) found that observers were significantly faster to indicate whether two probes placed along the contour were the same or different when they straddled a convex curvature extremum than when they straddled a locally identical concave extremum. This result demonstrates a part-superiority effect, i.e., a superior performance in making comparisons within a single part than across two parts. Importantly, this influence of sign of curvature disappeared in the case of the 'snake' shape. On this shape, the same local concavities no longer correspond to part boundaries, but rather to the concave sides of bends. Hence, Barenholtz and Feldman's results demonstrate that

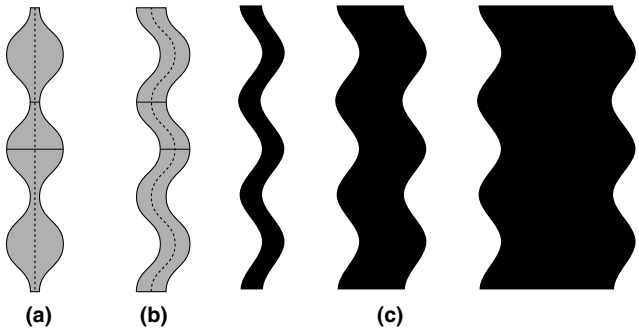

Fig. 4. Differences in shape arising from medial-axis geometry. In (a) the curvature of the bounding contour arises from local variations in the thickness of the shape around a straight axis. In (b) it arises from the curvature of the axis itself. (c) Stimuli used by Burbeck and Pizer (1995) in their study of the role of overall width in shape perception. With increasing shape width, observers perceive the axis of the shape as having systematically less curvature.

the role played by a concavity is modulated by nonlocal aspects of surface geometry—factors that determine its part decomposition.

2.3. Shape width

One by-product of the medial-axis representation is that it provides a perceptually natural measure of shape width (i.e., cross-axial thickness). Shape width has been shown to influence the perceived curvature of a shape's axis and the perceptual salience of its parts (Burbeck & Pizer, 1995; Siddiqi et al., 1999). Burbeck and Pizer estimated the perceived axis of a 'snake' shape (see Fig. 4(c)) using a shape-bisection task: observers indicated whether a briefly flashed dot appeared to the left or right of the *local* center. They found that as the width of a shape increases, its axis is perceived to be systematically less curved (having lower amplitude of modulation) even though the sinusoidal bounding contours are identical. Based on this and other results, Burbeck and Pizer proposed a model in which the bounding contours of a shape are detected with systematically weaker precision with increasing shape width. This entails that the perceptual salience of a local concavity or convexity on a shape—and hence of a part—decreases systematically with the overall width of the shape.

3. Surface geometry and visual completion

Although the influence of region-based geometry is well recognized in shape perception, as noted earlier, models of illusory and partially occluded contours typically take into account only the geometry of local inducing-contour pairs—in particular, the relative locations and orientations of the inducing edges.

Some recent work has pointed to the insufficiency of purely contour based mechanisms in visual completion. Yin, Kellman, and Shipley (2000) showed that surface color and texture can complement contour-based mechanisms of amodal completion, in that two surface fragments are more likely to group perceptually when they share a common color and texture. Recently, a more basic role for surface attributes has been demonstrated by Anderson, Singh, and Fleming (2002). They showed that factors involving the polarity of contrast at junctions, and the border ownership of contours, strongly modulate the visual completion of surfaces (see their "serrated edge illusion," Fig. 19). In particular, reversals in the sign of binocular disparity were shown to lead to altogether different visually completed surfaces—with the perceived surface lying on one or the other side of a visually interpolated contour—or, depending on contrast polarity, not being completed at all. Compelling examples of 3D completion displays (e.g., Gerbino, 2003; Tse, 1999a, 1999b; Van Lier & Wagemans, 1999) suggest that volumetric representations may also play a role in visual completion. These studies thus point to the insufficiency of purely contour-based approaches to visual completion.

In the context of amodal completion, two recent studies have presented direct evidence for the influence of sign of curvature. Liu, Jacobs, and Basri (1999) measured the strength of amodal grouping between pairs of inducing surface fragments using a depth-discrimination task. They found that amodal grouping was stronger when the interpolating contours were convex rather than concave. They argued for a model of amodal completion that predicts grouping strength by classifying pairs of inducing surface fragments into four types, depending on their convexity relationships (Fig. 5(a); see also Jacobs, 1996; Tse, 1999b).

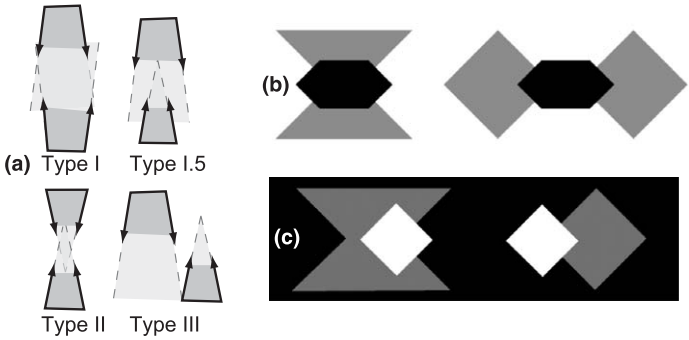

Fig. 5. (a) Liu et al.'s (1999) classification of the geometry of inducing surface fragments into four types, as a predictor of the strength of amodal completion. Type I is predicted to generate the strongest grouping; Type III the weakest. (b) and (c) Fantoni et al.'s (2005) experimental stimuli for estimating the perceived shape of partially occluded vertices.

Fantoni, Bertamini, and Gerbino (2005) estimated the amodally completed shapes of partly occluded corners using a vertex-separation task, and a probe-localization task. In the vertex-separation task, observers indicated whether the separation between two occluded vertices was greater or smaller than the length of a comparison line segment (see Fig. 5(b)). In the probe-localization task, observers indicated whether a briefly flashed line probe appeared inside or outside the amodally completed shape (see Fig. 5(c)). With both methods, Fantoni et al. found that the amodal contours were "more extrapolated"—i.e., closer to a corner—when the partly occluded surface was concave, than when it was convex. They proposed that a surface-based constraint involving the minimization of total area of the completed surface must supplement the contour-based determiners of amodal shape completion.

These results are significant in that they point to the insufficiency of contour geometry in predicting amodal completion—both the *strength* of grouping and the *shape* of interpolated contours. However, these studies have been performed only in the context of amodal completion. In this article, we extend this investigation to illusory contours. (As recent work has shown, modal and amodal completion do not necessarily generate the same shapes; see Anderson et al., 2002; Singh, 2004). More importantly, these studies have investigated only the role of sign of curvature. Consequently, although the proposed models capture the influence of local convexity, they do not explicitly take into account the influence of shape width or medial-axis geometry.

As noted in the previous section, the role played by a local convexity or concavity in shape perception is modulated strongly by region-based geometry. Our goal is to investigate the role of surface geometry more generally in determining illusory-contour shape, by examining interactions of the sign of curvature with shape width and medial-axis geometry. If the shapes of illusory contours are determined not simply by the sign of curvature, but also by nonlocal surface-based geometric factors—in a way that parallels the influence of these factors on perceived part structure—this would provide evidence for an intimate relationship between part-based representation of shape and visual mechanisms of shape completion. This, in turn, may provide a more general framework for the problem of shape interpolation.

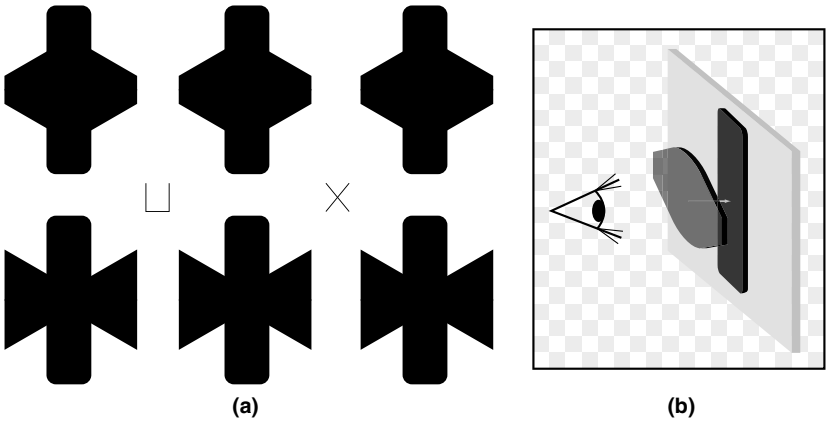

Fig. 6. (a) Stereoscopic illusory-contour configurations used in Experiment 1. (Cross-fusers should fuse the right two images; divergers the left two images.) When fused, the horizontally oriented shape (convex diamond or concave bowtie) is perceived as completed and floating in front of the vertical rectangle. These stimuli were used as the "standard" displays; observers adjusted the shape of a comparison figure (see Fig. 8) in order to match the shape of the illusory contours in the standard display. (b) Schematic depiction of the perceived depth layering.

4. Experiment 1: Sign of curvature and shape width

Experiment 1 uses illusory-contour displays that equate the local geometry of inducing-contour pairs, and manipulates the convexity of the enclosed surface. The displays contain two inducing surface fragments (on the sides of a vertically oriented rectangle; see Fig. 6), rather than three or four inducers, as in the standard Kanizsa-type displays. This configuration was employed because it forces the illusory contour to turn through a larger angle in proceeding from one inducing contour to the next. Smaller turning angles, as found in standard illusory-contour displays, allow less room for the shapes of the illusory contours to vary (see Fig. 7(a))—and thereby make it less likely that differences in the shapes of illusory contours can be detected (see Singh, 2004). Unlike the standard IC displays, however, this configuration requires binocular disparity to define the relative depth of the illusory contours, and to produce a vivid percept of their shape.

The two shape types used to manipulate the convexity of the enclosed surface are a *bowtie* with concavities that divide the shape into two distinct parts, and a convex *diamond*. Also manipulated are the width of the illusory surface and the separation between the inducing fragments. Illusory-contour shapes are measured using a parametric shape-adjustment method, in which observers adjust the shape of a comparison stimulus in order to match the shape of the illusory surface in the standard display (Singh, 2004).

4.1. Methods

4.1.1. Observers

Seventeen Rutgers University students participated in the experiment in exchange for course credit. All were naïve to the purpose of the experiment.

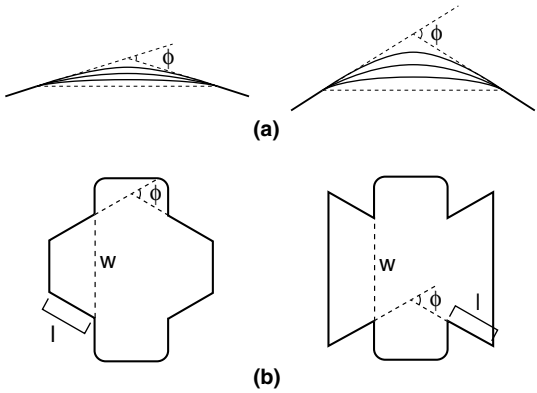

Fig. 7. (a) The turning angle is the total angle ϕ through which an interpolating contour must turn in proceeding from one inducing contour to the next. Note that smaller turning angles (left display), found in standard illusory-contour displays, allow less room for variation in the shape of the interpolating contour. (b) The parameters of shape width w, inducer length l, and turning angle ϕ are equated across the convex (diamond) and concave (bowtie) stimuli.

4.1.2. Stimuli and apparatus

Stimuli consisted of a stereoscopic illusory-contour display (the "standard") and an adjustable shape (the "comparison"). The standard display contained a horizontally oriented shape (bowtie or diamond) whose middle portion was camouflaged by an underlying vertically oriented rectangle (with rounded vertices) of the same luminance (0.03 cd/m^2; see Fig. 6(a)). The two overlapping shapes were presented against a gray background (7.8 cd/m^2).

The inducing fragments in the convex and concave displays were equated in terms of their local contour geometry. In both cases, the points of occlusion on the two inducing fragments on the lateral sides were horizontally aligned (prior to the introduction of binocular disparity), and the turning angle ϕ between each pair of inducing edges was 60° (see Fig. 7b). The visible length l of each inducing edge was 1.79°. The widths w of the inducing fragments, measured at the point of occlusion, were equated and could take one of two values: 2.21° or 4.42°. Straight vertical edges were added on both the sides, so that the shape width w could be manipulated independently of the inducer length l. The vertically oriented rectangle was 8.28° long and could take one of four widths, corresponding to four different values of inducer separation: 1.38°, 2.07°, 2.76°, and 3.45° of visual angle. The highest inducer separation was chosen such that, in the bowtie configuration, the two inducing edges on the same surface fragment (i.e., on one side of the vertical rectangle) do not intersect each other before intersecting the inducing edges from the opposite side. This ensures that, in the corresponding comparison shape, "zero smoothing" corresponds to a bowtie shape, and not to two separate triangles. The vertical rectangle was given a far disparity of 9.94 min of arc relative to the inducing surface fragments. This gave rise to the percept of a single illusory surface (bowtie or diamond) floating in front of the rectangle (see Fig. 6(b)).

The comparison display consisted of a single shape, whose lateral sides were identical to the visible inducing fragments of the illusory shape in the standard display (see Fig. 8). Observers adjusted the shape of its middle portion using a trackball, in order to match the shape of the illusory contours in the standard display. The adjusted shape was

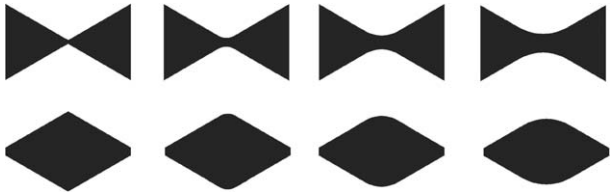

Fig. 8. Comparison shapes with different degrees of smoothing. Observers adjusted the smoothing level in order to match the perceived shape of the illusory contours in the standard display. The normalized smoothing index (see text for details) for the four levels of smoothing shown are 0, 0.273, 0.547, and 0.820, respectively.

parametrized in terms of a single variable—the smoothing level, defined by the size of 1-D box, or "averaging", convolution kernel needed to produce that level of smoothing (starting with the sharp vertices formed by the intersections of the linear extensions of the inducing edges). Smoothing was applied symmetrically to the upper and lower boundaries of the comparison shape. The larger the kernel size set by an observer, the greater the degree of smoothing. The range of possible kernel sizes was constrained so that the smoothing only affected the portion of the comparison display corresponding to the camouflaged portion of the shape in the standard display. The lateral sides of the comparison display, corresponding to the visible inducing edges in the standard display, were not affected. The kernel size set by the observer was normalized by the largest kernel size possible, thereby yielding a scale-invariant index of smoothing, ranging between 0 and 1 (see Singh, 2004).

The experiment was written in Matlab using the PsychToolbox libraries (Brainard, 1997; Pelli, 1997). The stimulus displays were presented on a high-resolution 22-in. monitor (Mitsubishi DiamondPro), under conditions of low ambient illumination. The stimuli were viewed from a distance of 126 cm through LCD shutter goggles (StereoGraphics CrystalEyes) that alternated at a rate of 150 Hz between the two eyes.

4.1.3. Design and procedure

The experimental design consisted of 2 inducer shapes (convex "diamond" and concave "bowtie") × 2 inducer widths × 4 inducer separations × 6 repetitions. Each observer thus performed 96 experimental shape adjustments. These were preceded by 16 practice adjustments.

Each trial began with a fixation cross in the center of the left half of the monitor followed by a standard display containing an illusory surface. After 5 s, a low-pitched tone signaled the onset of the comparison display in the right half of the monitor. Observers adjusted the shape of the comparison display using a trackball, under the constraints described above, in order to match the shape of the illusory contours in the standard display. Observers were allowed to take as much time as needed to make their adjustments.

4.2. Results and discussion

The graphs in Fig. 9 plot the data in terms of the normalized smoothing index. The graph on the left corresponds to the smaller shape width, the one on the right to the larger shape width.

Sign of curvature exerted a significant influence on illusory-contour shapes, $F(1,16) = 9.76$, $p < 0.01$. Illusory contours corresponding to convex shapes were perceived

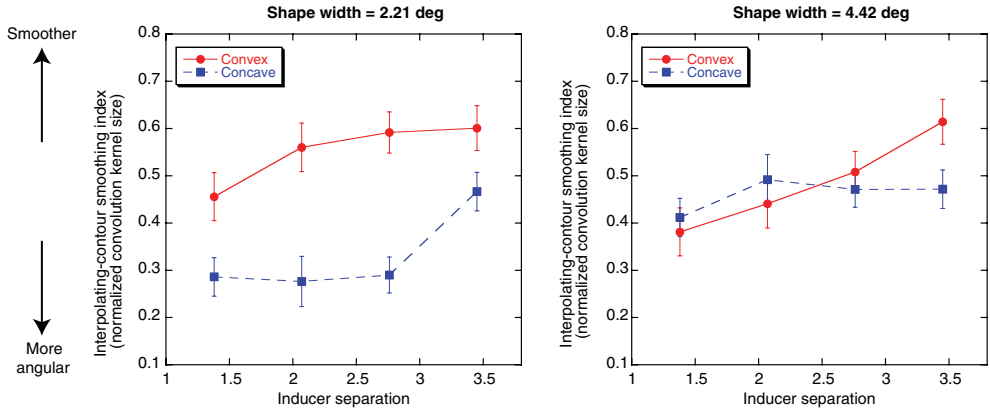

Fig. 9. Results of Experiment 1 showing the mean smoothing level set by the observers to match the shape of illusory contours. The results show a significant influence of inducer convexity on illusory-contour shape: Concave interpolating contours are systematically more angular than convex ones. Moreover, the influence of sign of curvature is substantially modulated by overall shape width: its influence is substantially weaker for wider shapes (graph on the right) than narrower ones (graph on the left). Error bars indicate standard errors of the mean.

as being more smoothed, whereas those corresponding to concave shapes were systematically more angular, i.e., closer to a corner. This effect of sign of curvature extends Fantoni et al.'s (2005) findings to the domain of illusory contours.

Importantly, the interaction between shape width and sign of curvature was also significant ($F(1,16) = 16.66$, $p < 0.001$), thereby showing that the role played by sign of curvature is modulated by shape width. As is clear from the plots in Fig. 9, sign of curvature exerts a much weaker influence for wider shapes (graph on the right) than it does for narrower ones (graph on the left). This result parallels the findings of Burbeck and Pizer (1995) in the context of shape bisection, where the influence of a local convexity or concavity was shown to weaken with increasing shape width. It is also consistent with Siddiqi et al.'s (1999) "part-protrusion axis" in their *shape triangle*—along which the same undulations of the bounding contour are perceived either as distinct parts, or as minor undulations, depending on the shape's width. The main effect of shape width was not statistically reliable ($F(1,16) = 1.73$, $p > 0.2$).

The main effect of inducer separation was significant ($F(3,48) = 8.77$, $p < 0.01$). As is evident from the plots, visually completed shapes became smoother with increasing size of the gap to be interpolated. Because the lengths of the inducing contours were kept fixed, this result may be interpreted as an influence of support ratio (i.e., the ratio of physically specified contour length to total contour length—physically specified plus interpolated; see Kellman & Shipley, 1991): the weaker this ratio, the smoother its shape (see also Fantoni & Gerbino, 2003). The lowest degree of smoothing observed is around 0.3; thus IC shapes are never perceived as being too close to the "corner" solution—i.e., linear extrapolants intersecting in a tangent discontinuity. In terms of Fantoni and Gerbino's (2003) model, this entails that the *good-continuation* or *GC* component (which models the extrapolation of each individual inducer) is insufficient to capture perceived IC shapes; there is a significant contribution of the *minimal-path* or *MP* component as well (which tends to minimize the total length of the interpolated contour).

The results of Experiment 1 demonstrate that the local geometry of inducing contours is not sufficient to predict the shape of illusory contours. In the stimuli used, local contour geometry was equated across the convex and concave shapes, and across the two shape widths used. Nevertheless, large differences were observed in illusory-contour shapes: convex contours were significantly more smoothed than concave ones, and the influence of sign of curvature was significantly modulated by the width of the shapes. Surface geometry thus plays an important role in determining the shape of illusory contours.

5. Experiment 2: Sign of curvature and medial-axis geometry

In Experiment 1, both shape types used had a straight axis and varying thickness around that axis. Experiment 2 adds a third configuration—the bent tube (see Fig. 10(a), left)—to the experimental design in order to manipulate medial geometry. As discussed above, in the bent tube, the curvature of the bounding contour arises entirely from the curvature of its axis; the thickness function of the shape itself is constant. In measuring the perceived shape of illusory contours in this display, however, there is a potential concern in using the same procedure as in Experiment 1. If the parametric adjustment of the comparison shape affects its two sides equally, as in Experiment 1, the convex and concave sides of the bent tube will be constrained to have identical shapes—and this would not permit a test of possible shape differences between the convex and concave sides of the bent tube.

Investigating the interaction between sign of curvature and medial-axis geometry requires that independent measurements be taken on the two sides of the illusory contours in the bent-tube configuration. However, there is an inherent difficulty in using parametric shape matching for this purpose. Even if observers are given the freedom to adjust the two

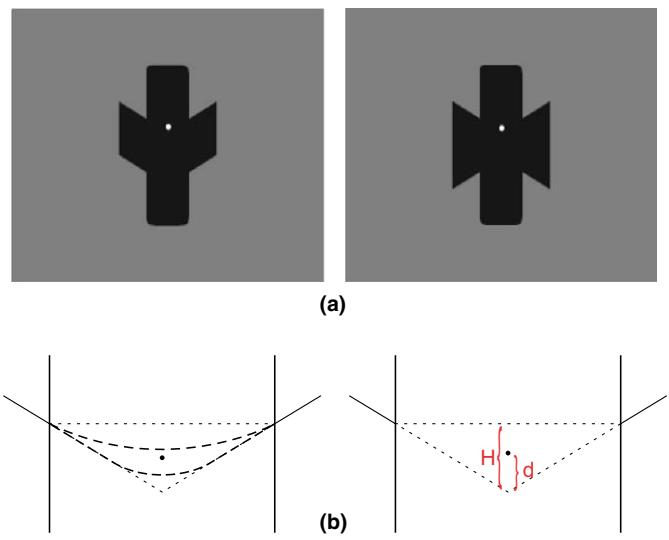

Fig. 10. (a) Monocular versions of stereoscopic stimuli used in Experiment 2. A dot flashes intermittently in the depth plane of the inducing contours. Observers adjusted its vertical position until it was perceived to lie on the illusory contour. (The dot is not drawn to scale here; it is shown larger for illustrative purposes.) (b) Observers' setting of dot position provides an estimate of the shape of the illusory contour. A normalized measure of smoothing is given by the ratio d/H: the larger this measure, the smoother the perceived illusory contour.

sides of the comparison shape separately, there is nevertheless a danger that their shape setting on one side may interfere with the adjustment of the other—e.g., via a bias toward keeping the two sides of the comparison shape symmetric or parallel. To circumvent this problem, we use a different experimental method to measure the shapes of illusory contours—one that does not rely on a separate comparison display for shape measurement (see Anderson & Barth, 1999). (Variants of this method, involving either the adjustment of a line probe (Fantoni et al., 2005) or dot-localization using a 2AFC method, (Guttman & Kellman, 2004) have also been employed.) A small dot flashes intermittently—its horizontal location constrained to be equidistant from the two points of occlusion. Observers' task is to adjust the vertical location of the dot so that it appears to lie *on* the boundary of the illusory surface. The intermittent flashing of the dot minimizes the influence that the presence of the dot may have on the perceived shape of the illusory contour. The measurement of the illusory-contour shape along this vertical cross-section allows for an estimation of its degree of smoothness: the closer the setting is to the intersection of the linear extensions of the two inducing edges, the more angular the illusory shape is (see Fig. 10(b)).

Two possible results may be predicted, depending on whether illusory-contour shape is determined simply by local convexity and concavity, or influenced by medial geometry as well. If it is determined simply by local sign of curvature, the shape difference between the convex and concave illusory contours should be the same, irrespective of whether the convexities and concavities belong to a shape with a straight axis and varying thickness (i.e., diamond or bowtie), or to one with a curved axis and constant thickness (i.e., bent tube). On the other hand, if illusory-contour shape is influenced by nonlocal geometry involving the medial axis, then the difference between convex and concave cases should be significantly smaller in the case of the bent tube. As noted earlier, the medial-axis geometry of bent tube is not consistent with a two-part structure, despite the presence of a concavity (see Siddiqi et al., 1999; Barenholtz & Feldman, 2003).

5.1. Methods

5.1.1. Observers

A new group of 15 undergraduate students from Rutgers University participated in the experiment in exchange for course credit. All were naïve to the purpose of the experiment.

5.1.2. Stimuli and apparatus

The primary difference from Experiment 1 was the absence of a comparison display. Only the stereoscopic illusory-contour display was presented. As before, the different inducer types were equated for the local geometry of the inducing-contour pairs—including turning angle (60°), visible inducing-contour length (1.79°), and inducing-surface width at the point of occlusion (2.76°). The vertical rectangle had a fixed length of 8.28°, and one of three possible widths: 1.38°, 2.42°, and 3.45°. The rectangle was given a far disparity of 9.94 min of arc relative to the inducing surface fragments. The viewing conditions and apparatus were the same as in Experiment 1.

5.1.3. Design and procedure

The factorial design was sign of curvature (convex, concave) × medial-axis configuration (straight-axis with varying thickness, curved axis with constant thickness) × 3 inducer

separations (1.38°, 2.42°, and 3.45°) × 2 measurement boundaries (upper, lower) × 5 repetitions. Each observer thus performed 120 experimental adjustments. These were preceded by 24 practice adjustments.

A stereoscopic illusory-contour display was presented in the center of the screen. Following a 1000 ms delay, a small light-gray dot (22.1 cd/m^2) of radius 2 pixels (1.66 min of arc) began to flash intermittently (*on* for 200 ms, *off* for 1200 ms). The dot was placed in the same depth plane as the inducing contours, and was constrained to lie either in the upper half of the illusory-contour display (for trials measuring the upper boundary) or in its lower half. The dot's horizontal position was constrained to be equidistant from the corresponding points of occlusion on the two sides. Its vertical position was to be adjusted by the observers. Observers' task was to set its vertical location using a trackball, so that it was perceived to lie *on* the boundary of the illusory surface.

Observers' settings were recorded in terms of the vertical distance d from the vertex formed by the intersection of the linear extensions of the two inducing contours (see Fig. 10(b)). This distance was then normalized by the maximal distance permissible—the height H of the triangle formed by the straight-line join of the two points of occlusion and the linear extensions of the two inducing contours. This normalized dot setting provides a scale-invariant measure of the degree of smoothing of the perceived illusory contour: the closer the setting is to 0, the more angular is the shape of the illusory contour.

At the beginning of each trial, the dot was initially located at the base of the triangle ($d = H$). Observers were thus allowed the option of positioning the dot near the base, even though this setting does not correspond to a legitimate smooth interpolant—but rather to a straight-line join of the two points of occlusion, generating tangent discontinuities at those points. (Apart from the fact that no theory of human visual completion predicts this "solution," it is noteworthy that even in Experiment 1, where the parametric shape adjustment permitted only smooth interpolating contours, observers' settings of contour smoothing remained well below the maximal smoothing allowed; mean settings of smoothing index never exceeded 0.75; see Fig. 9.) Because of the more time-consuming nature of the task used in Experiment 2—given the pauses between the brief flashes—this portion of the range was nevertheless included in order to provide an independent means of assessing whether or not each observer was performing the experimental task (the observers were undergraduate students, participating for course credit). Data from observers with more than a third of their settings within 5% of the distance to the base (i.e., $d/H > 0.95$) were excluded from the analysis. Data from two observers were excluded from the analysis based on this criterion. (The precise value of the criterion, in fact, makes little difference. The two observers had approximately half of their settings within a few pixels of the base, whereas the average corresponding proportion for the remaining observers was <1%.)

5.2. Results and discussion

The normalized data for the 13 observers are plotted in Fig. 11. As in Experiment 1, the main effect of convexity was highly significant ($F(1,12) = 47.64$, $p < 0.001$), with illusory contours being significantly more angular for locally concave contours than for locally convex ones. Importantly, the interaction between the sign of curvature and medial-axis geometry was also significant ($F(1,12) = 4.97$, $p < 0.05$). As is evident from the graphs in Fig. 11, the difference between the shapes of locally convex and locally concave contours is substantially greater for surface geometries involving a straight axis with a varying thickness

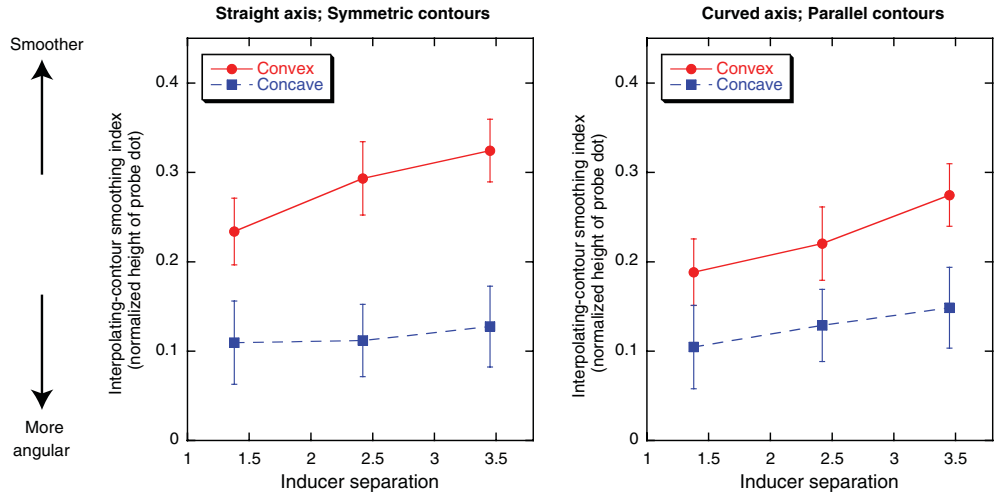

Fig. 11. Results of Experiment 2 showing the perceived smoothing index of illusory contours based on the normalized height of the probe dot set by the observers. As in Experiment 1, the results show a significant influence of inducer convexity on illusory-contour shape. Importantly, the medial geometry is also found to modulate the influence of sign of curvature. Error bars indicate standard errors of the mean. (Computation of within-observer variances showed the settings of individual observers were highly reliable. For most cells, standard deviations based on mean individual-observer variance were approximately four times smaller than those across observers.)

(concave bowtie versus convex diamond) than for those involving a curved axis with constant thickness (concave versus convex side of the bent tube). As noted above, this difference (and hence the interaction) is expected based on theories of surface-based shape representation. In the case of the diamond and bowtie shapes, the difference between convex and concave reflects a difference in part structure, whereas in the case of the bent-tube shape it does not—the local convexity and concavity arise from the curvature of the axis itself, rather than the presence of distinct parts. Thus, just as in the case of part-based representations of shape more generally (see, e.g., Barenholtz & Feldman, 2003; Siddiqi et al., 1999), the role that a local convexity or concavity plays in determining the shape of illusory contours is modulated by nonlocal aspects of surface shape involving medial-axis geometry.

Consistent with the results of Experiment 1, the main effect of inducer separation was also significant, $F(2, 24) = 25.12$, $p < 0.001$: illusory-contour shapes became systematically smoother with increasing separation between the inducing contours.

6. General discussion

Traditional approaches to image segmentation begin with the measurement of local edges and their integration, and interpolation, into extended contours. These contours are then used to delineate enclosed regions, and to propagate surface characteristics—"features" such as brightness or color—into these regions (e.g., Grossberg & Mingolla, 1985). This approach is sometimes taken to be supported by single-cell recordings, indicating that illusory contours are computed in the earliest visual cortical areas—primarily in V2 (Peterhans & von der Heydt, 1989; von der Heydt & Peterhans, 1989), with more recent

studies implicating V1 as well (Bakin, Nakayama, & Gilbert, 2000; Lee & Nguyen, 2001; Sheth, Sharma, Rao, & Sur, 1996).

Recent work in computational vision has adopted an alternative approach in which segmentation is informed right from the start by region-based attributes (e.g., Pao, Geiger, & Rubin, 1999; Sharon, Brandt, & Basri, 2000; Shi & Malik, 2000). This approach allows for a quick initial—though low-resolution—computation of salient regions in an image, which can guide high-resolution mechanisms of contour detection and integration. Such an approach is supported by recent fMRI studies on the human visual system, where higher cortical regions involved in surface processing, such as the LOC, have been shown to exhibit an extremely fast response (90–100 ms) to illusory-contour displays (Mendola, Dale, Fischl, Liu, & Tootell, 1999; Murray et al., 2002), thereby making the idea of feedback to V1 and V2 plausible. Moreover, these areas exhibit a response even when the crisply defined illusory contours have been destroyed by rounding off the pacmen inducers in Kanizsa displays (Stanley & Rubin, 2003), thereby indicating that the response to coarsely defined salient regions is not contingent upon the prior computation of contours.

Our results support the latter approach to image segmentation and perceptual organization, by demonstrating an influence of surface geometry on the shape of illusory contours. Models of visual completion—whether they be geometric, computational, or neural—have traditionally formulated the problem of shape completion as one of *contour* interpolation. Indeed, the geometry of visual shape interpolation has typically been encoded simply in terms of the local geometry of the inducing contours. The synthesis of illusory shapes, in particular, has been presumed to be driven by contour-based mechanisms—based on 1D constraints or interactions that apply along contours. Constraints involving surface geometry, e.g., the geometry of the region enclosed by these contours, have implicitly been assumed to not play a role.

The reported experiments manipulated a number of factors relating to surface geometry, while preserving the local geometry of inducing-contour pairs. In both experiments, the data exhibited large and systematic influences of surface geometry on the shape of illusory contours. First, locally concave inducers were found to generate systematically more angular illusory-contour shapes than corresponding locally convex inducers. This first result extends recent findings on amodal completion (Fantoni et al., 2005) to the domain of illusory contours. Importantly, however, our results also went beyond sign of curvature by demonstrating systematic influences of nonlocal surface geometry. The influence of sign of curvature on illusory-contour shape was found to be modulated significantly by the width of the shape and its medial geometry. Specifically, the shape difference between convex and concave illusory contours was systematically (i) greater for narrower shapes than wider ones, and (ii) greater for shapes whose surface geometry is consistent with a two-part structure than for those with a curved axis and constant thickness (where the local concavity does not correspond to a part boundary). The observed shape difference between the convex and concave contours in the curved-axis condition is especially pertinent, since it suggests that a constraint based on minimizing total surface area (e.g., Fantoni et al., 2005) or on the attenuation of color diffusion with distance from a contour (Grossberg, 1987; Grossberg & Mingolla, 1985) is not sufficient by itself to explain the influence of surface geometry on illusory-contour shape. Although an area-based measure can explain the observed difference in illusory-contour shape between the diamond and bowtie configurations used here, it cannot account for the shape difference observed between the convex

and concave contours in the bent-tube configuration. The area-based support ratio, for instance, would of course be the same for the two sides of the bent tube.

The influence of sign of curvature, shape width, and medial geometry on illusory-contour shape parallels the role that these factors play in visual shape representation more generally—in particular, in organizing shapes into parts (Barenholtz & Feldman, 2003; Burbeck & Pizer, 1995; Siddiqi et al., 1999; Singh & Hoffman, 2001). Geometric factors that increase the likelihood and salience of a perceptual part boundary (such as the presence of a local concavity, a narrower overall width, and medial geometry being consistent with a multi-part structure) also result in more angular interpolated shapes. These findings suggest that the influence of these factors on the shape of illusory contours is intimately related to (and possibly mediated by) the role they play in visual part decomposition. A detailed investigation of this relationship remains an important topic for future research.

7. Conclusion

Even at the level of illusory "contours," there is an early contribution of region-based geometry which is sensitive to nonlocal shape properties involving medial-axis representation and part decomposition. Models of illusory-contour synthesis must incorporate the role of region-based geometric factors in a way that parallels their role in organizing visual shape representation more generally.

Acknowledgements

We are grateful to Bart Anderson, Jacob Feldman, Eileen Kowler, and Larry Maloney for helpful discussions, and Marco Bertamini and a second anonymous referee, as well as the editors of the special issue, for their comments and suggestions on the manuscript. This work was supported by NSF grant BCS-0216944.

References

Anderson, B. L., & Barth, H. C. (1999). Motion-based mechanisms of illusory contour synthesis. *Neuron, 24*, 433–441.

Anderson, B. L., Singh, M., & Fleming, R. (2002). The interpolation of object and surface structure. *Cognitive Psychology, 44*, 148–190.

Attneave, F. (1971). Multistability in perception. *Scientific American, 225*, 63–71.

Bakin, J., Nakayama, K., & Gilbert, C. D. (2000). Visual responses in monkey areas V1 and V2 to three-dimensional surface configurations. *The Journal of Neuroscience, 20*, 8188–8198.

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.

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). Obligatory edge assignment in vision: The role of figure and part segmentation in symmetry detection. *Journal of Experimental Psychology: Human Perception and Performance, 21*, 1323–1342.

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

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

Bertamini, M., & Hulleman, J. (2006). Amodal completion and visual holes (static and moving). *Acta Psychologica, 123*, 55–72.

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

Brady, M., & Asada, H. (1984). Smoothed local symmetries and their implementation. *The International Journal of Robotics Research, 3*, 36–61.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10*, 433–436.

Braunstein, M. L., Hoffman, D. D., & Saidpour, A. (1989). Parts of visual objects: An experimental test of the minima rule. *Perception, 18*, 817–826.

Burbeck, C. A., & Pizer, S. M. (1995). Object representation by cores: Identifying and representing primitive spatial regions. *Vision Research, 35*(13), 1917–1930.

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*, 313–321.

De Winter, J., & Wagemans, J. (2004). Contour-based object identification and segmentation: Stimuli, norms and data, and software tools. *Behavior Research Methods Instruments and Computers, 36*, 604–624.

De Winter, J., & Wagemans, J. (2006). Segmentation of object outlines into parts: A large-scale integrative study. *Cognition*.

De Wit, T., & Van Lier, R. (2002). Global visual completion of quasi-regular shapes. *Perception, 31*, 969–984.

Elder, J. H., & Goldberg, R. M. (2002). Ecological statistics of Gestalt laws for the perceptual organization of contours. *Journal of Vision, 2*(4), 324–353.

Fantoni, C., Bertamini, M., & Gerbino, W. (2005). Contour curvature polarity and surface interpolation. *Vision Research, 45*, 1047–1062.

Fantoni, C., & Gerbino, W. (2003). Contour interpolation by vector-field combination. *Journal of Vision, 3*(4), 281–303.

Feldman, J., & Singh, M. (2005). Information along contours and object boundaries. *Psychological Review, 112*(1), 243–252.

Field, D., Hayes, A., & Hess, R. (1993). Contour integration by the human visual system: Evidence for a local "association field". *Vision Research, 33*(2), 173–193.

Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research, 41*(6), 711–724.

Gerbino, W. (2003). The joint. *Acta Psychologica, 114*, 331–353.

Grossberg, S. (1987). Cortical dynamics of three-dimensional form, color, and brightness perception: II. Binocular theory. *Perception & Psychophysics, 41*(2), 117–158.

Grossberg, S., & Mingolla, E. (1985). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review, 92*(2), 173–211.

Guttman, S. E., & Kellman, P. J. (2004). Contour interpolation revealed by a dot localization paradigm. *Vision Research, 44*, 1799–1815.

Heitger, F., von der Heydt, R., Peterhans, E., Rosenthaler, L., & Kübler, O. (1998). Simulation of neural contour mechanisms: Representing anomalous contours. *Image and Vision Computing, 16*, 407–421.

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.

Horn, B. K. P. (1983). The curve of least energy. *ACM Transactions on Mathematical Software, 9*(4), 441–460.

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

Jacobs, D. (1996). Robust and efficient detection of salient convex groups. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 18*, 23–37.

Kellman, P., & Shipley, T. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology, 23*, 141–221.

Kellman, P., Guttman, S., & Wickens, T. (2001). Geometric and neural models of object perception. In T. Shipley & P. Kellman (Eds.), *From fragments to objects: Segmentation and grouping in vision* (pp. 183–245). Oxford, UK: Elsevier Science.

Kimia, B., Frankel, I., & Popescu, A. (2003). Euler spiral for shape completion. *International Journal of Computer Vision, 54*(1/2), 157–180.

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

Koenderink, J., & van Doorn, A. (1982). The shape of smooth objects and the way contours end. *Perception, 11*, 129–137.

Koffka, K. (1935). *Principles of Gestalt Psychology*. London: Lund Humphries.

Lamote, C., & Wagemans, J. (1999). Rapid integration of contour fragments: From simple filling-in to parts-based shape description. *Visual Cognition, 6*(3/4), 345–361.

Lee, T., & Nguyen, M. (2001). Dynamics of subjective contour formation in the early visual cortex. *Proceedings of the National Academy of Sciences, USA, 98*, 1907–1911.

Leyton, M. (1988). A process-grammar for shape. *Artificial Intelligence, 34*(2), 213–247.

Liu, Z., Jacobs, D., & Basri, R. (1999). The role of convexity in perceptual completion: Beyond good continuation. *Vision Research, 39*, 4244–4257.

Mendola, J., Dale, A., Fischl, B., Liu, A., & Tootell, R. (1999). The representation of illusory and real contour in cortical visual areas revealed by fMRI. *Journal of Neuroscience, 19*, 8560–8572.

Michotte, A. (1950). On phenomenal permanence: facts and theories. *Acta Psychologica, 7*, 293–322.

Michotte, A., Thinès, G., & Crabbé, G. (1991). Amodal completion of perceptual structures. In G. Thinés, A. Costall, & G. Butterworth (Eds.), *Michotte's experimental phenomenology of perception* (pp. 140–167). Hillsdale, NJ: Erlbaum (Originally published in French 1964).

Mumford, D. (1994). Elastica and computer vision. In C. L. Bajaj (Ed.), *Algebraic geometry and its applications* (pp. 491–506). New York: Springer-Verlag.

Murray, M., Wylie, G., Higgins, B., Javitt, D., Schroeder, C., & Foxe, J. (2002). The spatiotemporal dynamics of illusory-contour processing: Combined high-density electrical mapping, source analysis, and functional magnetic resonance imaging. *Journal of Neuroscience, 22*, 5055–5073.

Nakayama, K., He, Z., & Shimojo, S. (1995). Visual surface representation: A critical link between lower-level and higher-level vision. In S. M. Kosslyn & D. N. Osherson (Eds.). *Visual Cognition: An invitation to Cognitive Science (2nd ed.)* (Vol. 2, pp. 491–506). Cambridge, MA: MIT Press.

Neumann, H., & Mingolla, E. (2001). Computational neural models of spatial integration in perceptual grouping. In T. Shipley & P. Kellman (Eds.), *From fragments to objects: Segmentation and grouping in vision* (pp. 353–400). Oxford, UK: Elsevier Science.

Pao, H., Geiger, D., & Rubin, N. (1999). Measuring convexity for figure-ground segmentation. *Computer Vision and Pattern Recognition, Proceedings*, 948–955.

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

Pelli, D. G. (1997). The Videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*, 437–442.

Peterhans, E., & von der Heydt, R. (1989). Mechanisms of contour perception in monkey visual cortex ii. Contours bridging gaps. *Journal of Neuroscience, 9*, 1749–1763.

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

Sekuler, A. B., Palmer, S. E., & Flynn, C. (1994). Local and global processes in visual completion. *Psychological Science, 5*, 260–267.

Sharon, E., Brandt, A., & Basri, R. (2000). Fast multiscale image segmentation. *Computer Vision and Pattern Recognition, Proceedings*, 70–77.

Sheth, B. R., Sharma, J., Rao, C., & Sur, M. (1996). Orientation maps of subjective contours in visual cortex. *Nature, 274*, 2110–2115.

Shi, J., & Malik, J. (2000). Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 22*, 888–905.

Siddiqi, K., Kimia, B., Tannenbaum, A., & Zucker, S. W. (2001). On the psychophysics of the shape triangle. *Vision Research, 41*(9), 1153–1178.

Siddiqi, K., Kimia, B. B., Tannenbaum, A., & Zucker, S. W. (1999). Shape, shocks and wiggles. *Image and Vision Computing, 17*, 365–373.

Siddiqi, K., Tresness, K. J., & Kimia, B. B. (1996). Parts of visual form: Psychophysical aspects. *Perception, 25*, 399–424.

Singh, M. (2004). Modal and amodal completion generate different shapes. *Psychological Science, 15*, 454–459.

Singh, M., & Fulvio, J. M. (2005). Visual extrapolation of contour geometry. *Proceedings of the National Academy of Sciences, USA, 102*(3), 939–944.

Singh, M., & Hoffman, D. D. (1999). Completing visual contours: The relationship between relatability and minimizing inflections. *Perception & Psychophysics, 61*, 636–660.

Singh, M., & Hoffman, D. D. (2001). Part-based representations of visual shape and implications for visual cognition. In T. Shipley & P. Kellman (Eds.), *From fragments to objects: Segmentation and grouping in vision. Advances in Psychology* (Vol. 130, pp. 401–459). Oxford, UK: Elsevier Science.

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

Stanley, D. A., & Rubin, N. (2003). fMRI activation in response to illusory contours and salient regions in the human lateral occipital complex. *Neuron, 37*, 323–331.

Tse, P. U. (1999a). Volume completion. *Cognitive Psychology, 39*, 37–68.

Tse, P. U. (1999b). Complete mergeability and amodal completion. *Acta Psychologica, 102*, 165–201.

Ullman, S. (1976). Filling-in the gaps: The shape of subjective contours and a model for their generation. *Biological Cybernetics, 25*, 1–6.

Van Lier, R. (1999). Investigating global effects in visual occlusion: From a partly occluded square to the back of a tree trunk. *Acta Psychologica, 102*, 203–220.

Van Lier, R., de Wit, T. C. J., & Koning, A. (2006). Con-fusing contours & pieces of glass. *Acta Psychologica, 123*, 41–54.

Van Lier, R., Van der Helm, P., & Leeuwenberg, E. (1994). Integrating global and local aspects of visual occlusion. *Perception, 23*, 883–903.

Van Lier, R., Van der Helm, P., & Leeuwenberg, E. (1995). Competing global and local completions in visual occlusion. *Journal of Experimental Psychology: Human Perception and Performance, 21*, 571–583.

Van Lier, R., & Wagemans, J. (1999). From images to objects: Global and local completions of self-occluded parts. *Journal of Experimental Psychology: Human Perception and Performance, 25*, 1721–1741.

von der Heydt, R., & Peterhans, E. (1989). Mechanisms of contour perception in monkey visual cortex. I. Lines of pattern discontinuity. *Journal of Neuroscience, 9*, 1731–1748.

Williams, L. R., & Jacobs, D. W. (1997). Stochastic completion fields: A neural model of illusory contour shape and salience. *Neural Computation, 9*, 837–859.

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

Yin, C., Kellman, P., & Shipley, T. (2000). Surface integration influences depth discrimination. *Vision Research, 40*, 1969–1978.