

Geometric determinants of shape segmentation: Tests using segment identification

Elias H. Cohen ^{*}, Manish Singh

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

Received 16 February 2007; received in revised form 6 June 2007

Abstract

The geometric determinants of shape decomposition were studied using a performance-based method. Observers' identification of contour segments was shown to be systematically modulated by their curvature properties, and by the geometric properties of the enclosed region. Specifically, negative minima of contour curvature provided the best segment boundaries. Segments with negative-minima boundaries were identified with greater accuracy than those with positive maxima or inflection boundaries of comparable length. Additionally, segment identification was shown to be determined by contour length, the turning angle at part boundaries, and the width at the part's base (hence the part's protrusion). The results indicate that part decomposition is an automatic process. Moreover, this process is graded, i.e. parts are more strongly segmented, or more likely to be perceived, according to the strength of many geometric determinants.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Shape representation; Part segmentation; Parts; Contour curvature; Structural descriptions; Part salience

1. Introduction

Understanding of the neural basis of visual shape has, for many years, been largely restricted to very early or very late stages of processing. Early visual cortical areas (such as area V1) are known to respond to simple local image features, such as oriented edges (Hubel & Wiesel, 1959, 1962). Work on late visual processing (such as area IT), on the other hand, has shown selectivity for high-level categories of complex recognizable shapes, such as faces or hands (Desimone, Albright, Gross, & Bruce, 1984; Gross, Rocha-Miranda, & Bender, 1972; Perrett, Rolls, & Caan, 1982; Tanaka, Saito, Fukada, & Moriya, 1991). However, processing at intermediate levels, responsible for transforming local image measurements into global representations of objects, is only beginning to be understood

(Pasupathy & Connor, 1999, 2001, 2002). It is at these intermediate-level representations of shape that we focus our psychophysical investigation. Despite a great deal of discussion of parts in the shape literature, surprisingly few performance-based methods have been developed to investigate the part-based nature of visual shape. This paper reports a series of psychophysical experiments designed to quantify geometric influences on the part-based representation of visual shape. In doing so, it provides some indication of how the gap between local measurements of edges and the global representation of shape might be bridged.

1.1. Shape and part

A great deal of information about an object's shape is carried by its *occluding contour* (Attneave, 1954; Koenderink, 1984). Indeed, recognition performance with silhouettes has sometimes been found to be almost as good as with shaded drawings of 3D objects (Hayward, Tarr, &

^{*} Corresponding author. Address: Department of Vision Sciences, State University of New York, State College of Optometry, 33 West 42nd Street, New York, NY 10036, USA.

E-mail address: ecohen@sunyopt.edu (E.H. Cohen).

Corderoy, 1999). Recent work using single-cell recordings in area V4 in the primate visual cortex has investigated the representation of this important source of information, and found evidence for a *piecewise* coding of shape. Specifically, Pasupathy and Connor (1999) demonstrated a selectivity of V4 cells to contour curvature as determined by *curvature polarity* (whether a contour segment has positive or negative curvature—i.e., bounds a locally convex or concave region) and *curvature magnitude* (or contour “acuteness”). Different populations of cells were found to respond selectively to either convex or concave segments of an occluding contour. In subsequent studies employing whole shapes, Pasupathy and Connor (2001, 2002) found evidence for a distributed coding of shape, with different units within a population responding to either convex or concave segments at different angular positions on the shape (relative to its center). The representation of shape in V4 thus appears to be a *piecewise* representation, in that different units in V4 encode different—either convex or concave—pieces or segments of a shape’s occluding contour.

The notion of a part-based representation has also played an important role in object recognition and behavioral approaches to visual shape. Although any subset of a shape may, in a sense, be considered its “part,” the perceptual notion of a part corresponds to those subsets of a shape that are naturally perceived as being semi-independent units—hence perceptually separable from the rest of that shape. Parts are thus more than the arbitrary pooling of earlier stage inputs; they are rather the perceptually and psychologically salient visual units of a shape’s representation. Organizing shape representations in terms of parts—with smaller parts hierarchically nested within larger parts—allows one to separate the representation of the shape of each individual part from the representation of the spatial relationships between the parts. This, in turn, leads to a more robust representation of shape—one that is more stable across viewing conditions—e.g., changes in the articulated pose of an object, or an observer’s vantage point with respect to it (Biederman, 1987; Hoffman & Richards, 1984; Marr & Nishihara, 1978; Palmer, 1977). Furthermore, hierarchical structure in shape representation is a central aspect of visual experience. The ability to direct action to different levels of object structure is crucial to visually guided manipulation and interaction. People, arms, hands, and fingers are all natural candidates for visual attention, semantic labeling, and motor interaction. Understanding the geometric determinants of shape segmentation thus allows one to understand how these separable units of shape are perceptually generated.

Part-based representation of shape, which central to many theories of high-level vision and object recognition, has generally lacked systematic psychophysical investigation. Compelling psychophysical support for part-based representation of shape should ideally take the following form: (i) demonstration that human vision automatically divides complex shapes into smaller units, (ii) demonstra-

tion that there are consistent and predictable rules that dictate the segmentation of shape into these sub-units, and that these sub-units determine the storage of visual information, and finally, (iii) demonstration that these same sub-units dictate the extraction of other visual properties, such as orientation, location, and size. The following investigations were undertaken with points (i) and (ii) as the guiding framework. (For work investigating point (iii), see Cohen & Singh, 2006; Denisova, Singh, & Kowler, 2006).

Parts have generally been studied from the point of view of high-level vision, as categorical units of object representation (Tversky & Hemenway, 1984) and units important for recognition and naming (e.g., Biederman, 1987; Biederman & Cooper, 1991). However, what a part *is* from the point of view of bottom-up visual processing (i.e., low-level mechanisms of visual segmentation) is less clear. Our goal is to characterize the notion of a visual part in concrete psychophysical terms, based on observers’ performance in an objectively-defined task. Unlike many previous studies, we use unfamiliar randomly-generated shapes, in order to focus on the geometric properties of the bounding contour. In a part-based account, certain portions of a shape constitute natural units of representation for the visual system, and therefore should be much more readily identifiable than other portions of comparable size. We therefore use segment identification as the operational test for naturalness of shape parts. The “null hypothesis” then becomes identification performance as might be predicted by a decomposition-free account of shape representation (such as one based on unstructured templates). Under such an hypothesis, any portion of a shape should be equally easy or difficult to identify, as long as its size is preserved.

In this paper, we focus on one source of shape information, namely occluding contour. Specifically, we examine the geometric properties of a contour segment that make it more or less identifiable. Experiments 1 and 2 compare contour curvature landmarks that potentially provide natural *boundary cues* for parts. Experiments 3 and 4 examine geometric factors that determine the *salience* of a part’s representation.

1.2. Contour curvature and information content

Attneave (1954) observed that objects’ boundaries—their occluding contours in the projected image—have high information content as they signify the greatest change in image characteristics. Koenderink (1984) has shown, more specifically, that the occluding contour carries a great deal of information about 3D shape: for smooth shapes, the sign of curvature of the occluding contour directly informs one of the sign of Gaussian curvature of the 3D shape.

Similarly, along occluding contours, Attneave (1954) observed that information is concentrated at extrema of contour curvature—points where the change signified by curvature is the highest (i.e. points where the magnitude

of curvature is locally maximal; see Norman, Phillips, & Ross, 2001, for quantification of visual sensitivity to curvature extrema, and De Winter & Wagemans, 2004, for an available database of observer responses). Although Attneave did not differentiate between different types of extrema based on the sign or polarity of curvature, subsequent work has demonstrated that extrema may be divided into multiple categories, including positive maxima and negative minima of curvature (see Fig. 1; Koenderink & Van Doorn, 1982; Koenderink, 1984; Hoffman & Richards, 1984; Richards & Hoffman, 1985; Leyton, 1989). Curvature extrema have also been shown to provide important perceptual landmarks on surfaces (Phillips, Todd, Koenderink, & Kappers, 2003). A more recent informational analysis has shown that negative minima (local peaks of concave curvature) carry greater information than positive maxima (local peaks of convex curvature; Feldman & Singh, 2005). The differential role of positive maxima and negative minima will play a central role in the experiments reported in this paper.

Their status as information peaks has made extrema of contour curvature the subject of much work in shape representation. It is theorized that curvature extrema play a significant role in the visual representation of shape—either locally, by receiving heightened emphasis in local contour representation, or globally, through their importance as landmarks utilized by mechanisms of shape segmentation and encoding (Barenholtz, Cohen, Feldman, & Singh, 2003; Cohen, Barenholtz, Singh, & Feldman, 2005; Bertamini & Farrant, 2005).

The latter hypothesis derives from the *minima rule* (Hoffman & Richards, 1984), which postulates that the visual system uses negative minima of curvature to define candidate boundaries between parts. Hoffman and Richards proposed this rule based on the principle of transversality, namely, that two randomly interpenetrating surfaces generically create a concave discontinuity at the locus of their intersection. It should be noted that the minima rule also embodies a fundamental asymmetry between positive and negative curvature: only curvature extrema in concave regions (i.e., negative minima) are postulated to serve as candidate points for shape segmentation; equivalent curva-

ture extrema in convex regions (i.e., positive maxima) do not have this status.

1.3. Contour integration and grouping

A parallel motivation for the role of curvature extrema in shape segmentation derives from work on contour integration. It has been demonstrated that the turning angles between successive local elements (the discrete version of curvature) in a discretely-sampled contour play a critical role in visually integrating these elements into the representation of a single extended contour (Field, Hayes, & Hess, 1993; Feldman, 1997; Pettet, McKee, & Grzywacz, 1998; Geisler, Perry, Super, & Gallogly, 2001). Contour integration, or grouping, is strongest for the smallest turning angles, and weakens with their increasing magnitude. The greater the turning angle between two given contour fragments, therefore, the less likely it is that they will group perceptually into a single unit. If the representation of smooth contours is governed by a similar dependence, curvature extrema would constitute the “weakest links” along a contour—thereby suggesting that they provide the most natural break points for segmenting a contour.

Studies of contour integration do not differentiate between curvature of positive and negative sign, because they generally employ open contours for which the sign of curvature is not canonically defined. However, in the related domain of visual completion of partially-occluded shapes, Liu, Jacobs, and Basri (1999) demonstrated an added influence of convexity on the amodal grouping of contours. Concave pairs of contour fragments were less likely to be perceptually grouped into the representation of a single surface than corresponding convex pairs. This study thus suggests that negative minima, even more than positive maxima, may serve as points of weak connectedness—giving further strength to the idea that they function as natural segmentation points for the visual system (a similar asymmetry in grouping between positive maxima and negative minima has been shown by Singh & Hoffman, 1998, in the context of perceptual transparency).

1.4. Shape segmentation and curvature extrema

Despite a strong theoretical basis and compelling visual examples, direct performance-based psychophysical evidence for automatic shape segmentation at negative minima is currently limited. In recent years, behavioral work has begun to show *differential representation* of curvature extrema based on their sign of contour curvature. Barenholtz et al. (2003) and Cohen et al. (2005) showed that observers are much more sensitive to shape changes involving negative minima of curvature than those involving positive maxima. Bertamini (2001) showed that observers are faster and more accurate in making positional (height) comparisons between an external landmark and a convex extremum, vs. a concave extremum. A third study by




Fig. 1. The figure demonstrates three types of curvature landmarks, positive maxima, negative minima, and inflections.

Lamote and Wagemans (1999), involving the detection of small gaps in closed contours, showed that observers were better able to detect gaps that fell on positive maxima of curvature, relative to those that fell on negative minima. These studies point to the distinct roles played by the two types of curvature extrema in various tasks. Rather than suggesting that one type of extremum is emphasized more or less in visual representation than the other, these studies taken together (along with physiological work in area V4; see Pasupathy & Connor, 1999, 2001) provide evidence for the hypothesis that negative minima and positive maxima are represented differentially.

But what role does the differential status of these two types of curvature extrema play in the visual representation of shape? Understanding of this issue is still quite incomplete. Behavioral evidence for part segmentation based on contour geometry thus far has involved one of two difficulties. Performance-based measures have mostly involved relatively indirect inference of part segmentation from other visual phenomena, such as judgment of object symmetry (Baylis & Driver, 1994), attentional superiority within parts (Barenholtz & Feldman, 2003; Vecera, Behrmann, & Filapek, 2001; Watson & Kramer, 1999), or visual search for shapes (Hulleman, te Winkel, & Boselie, 2000; Wolfe & Bennett, 1997; Xu & Singh, 2002). In an unpublished study described briefly in a review paper (Baylis & Driver, 1995), observers viewed stimuli composed of a vertical rectangle divided by a single jagged contour into two lateral regions, one red and one green. Observers were instructed to memorize either the red or green object. In the test phase, a horizontal section of the original stimulus or a similar distracter section was presented. Observers were asked to indicate whether this section matched the initial stimulus. Results showed that observers were faster and more accurate to verify the presence of the test section if it contained a contour segment that was largely convex (with regard to the assigned study color) as opposed to primarily concave. This study brings to light a question that remains open throughout studies of part decomposition. A tendency to favor convex stimuli is well documented, dating back to Rubin's observations (Rubin, 1921). Since parts defined by negative minima will often be largely—but not entirely—convex,¹ it has been impossible based upon previous research to determine which factor is dominant in visual part segmentation: Are parts defined simply through convexity—either contour-based (Latecki & Lakamper, 1999; Vaina & Zlateva, 1990) or region-based (Rosin, 2000)—or are they segmented explicitly at negative minima?

At the other end of the spectrum are experimental studies where observers are asked to explicitly draw part cuts on line drawings of familiar objects by hand (e.g. De Winter & Wagemans, 2004, 2006; Siddiqi, Tresness, & Kimia,

1996). De Winter and Wagemans (2004, 2006) have recently conducted several large scale studies which analyzed the hand-drawn part cuts generated by hundreds of observers. In doing so, they were able to quantify several geometric influences on part segmentation. This included demonstrating that, when asked to draw cuts, observers generally selected negative minima as segmentation points more often than other locations (although they also selected inflections and positive maxima). De Winter and Wagemans' studies also demonstrated that top-down factors, such as the recognizability of shapes as known objects, interact with bottom-up geometric factors. Specifically, cognitive knowledge of an object's functional parts mediates the number and placement of part cuts drawn by observers. For instance, one's knowledge that a wing is a functional part of a bird may lead one to segment it, regardless of the specific geometry of the occluding contour. The reliance on familiar shapes means, however, that contour geometry cannot be manipulated parametrically—thus making it difficult to quantify the influence of contour geometry in the absence of recognition. Moreover, these studies do not provide evidence for the *automaticity* of part segmentation.

One study using 3D surfaces of revolution, shown as structure-from-motion displays, tested observers' memory for sections of the studied shape (Braunstein, Hoffman, & Saidpour, 1989). When observers were asked to choose between four alternatives—a part with negative minima endpoints, a similar part with positive-maxima endpoints, and two distractors—observers selected the negative-minima option roughly 66% of the time. However, when the test options included only one veridical segment, observers were just as likely to select maxima-defined segments as minima-defined segments. Clearly, more definitive evidence is needed. The current work examines whether there is an automatic advantage for geometrically determined segment classes. Additionally, we will expand this investigation further than simply investigating which *type* of part is preferred, to examine a wider class of geometric influences on part segmentation.

The current experiments use a performance-based segmentation method that taps more directly into visual mechanisms of part segmentation. We employ two versions of this method: segment verification (Experiments 1 and 2) and segment identification (Experiments 3 and 4). The method allows a quantitative measure of the perceptual naturalness of candidate parts defined by specific geometric attributes. We are thus able to directly compare the perceptual efficacy of different extrema types (Experiment 1 and 2) on a wide class of randomly-generated shapes. Additionally, we compare the perceptual naturalness of segments defined by uniform sign of curvature (delineated by inflection points—where contour curvature switches its sign) to those bounded by positive maxima and negative minima (Experiment 2). Finally, in addition to examining geometric determinants of perceived part boundaries, it is important to understand geometric deter-

¹ Negative minima, by definition, lie in concave regions of the bounding contour. Hence parts segmented at smooth negative minima will invariably contain sections of negative curvature.

minants of part salience. We examine the influence of geometric properties on *strength* of segmentation—that determine a graded spectrum of shape decomposition (Experiments 3 and 4).

2. Experiments

Experiments 1 and 2 employed a segment-verification task (SVT). Observers were shown a test segment, which was followed by a mask, then an entire probe shape (see Fig. 2). Test segments matched a portion of the probe shape on 50% of trials. The observer's task was to indicate whether or not each trial contained a match between the test segment and the probe shape.

The SVT relies on a similar rationale as a task employed by Palmer (1977) to demonstrate the hierarchical structure of perceptual representation. In Palmer's task, observers were asked to verify the presence of a specific configuration of straight-line segments within a larger figure made up of such segments. Palmer's data showed that observers were fastest to identify a group of segments that fell within the same perceived hierarchical unit. The SVT used here, similar to that employed by Waeytens, Hanouille, Wagemans, and d'Ydewalle (1994) (see also Leek, Reppa, & Arguin, 2005), differs in several respects. Stimuli used in these experiments were randomly-generated smooth outline shapes. We were thus able to explore the process of segmentation over a wide class of smooth 2D shapes, and test which geometric properties of the occluding contour are most predictive of part segmentation. Half of the trials presented segments selected from the test shape. In mismatch trials, non-target segments were generated via a random distortion of a segment taken from the shape (described below). This method allowed multiple levels of task difficulty (created by manipulating the magnitude of vertex displacement) and allowed measurement of discrimination performance using d' .

The presentation order—test segment before the probe shape—was chosen based on observers' difficulty with the task in pilot studies when the order was reversed (yielding

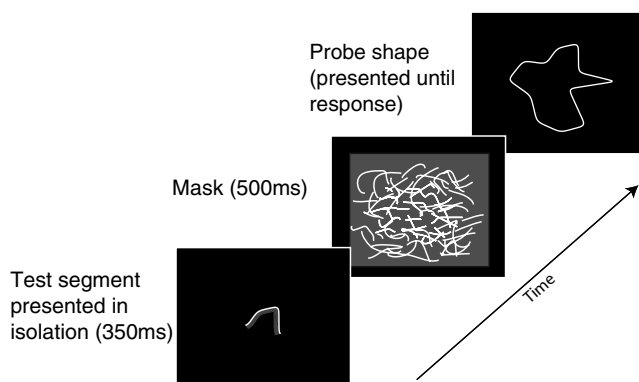


Fig. 2. An illustration of the trial sequence used in Experiments 1 and 2.




Fig. 3. Test segment types used in Experiments 1a and 1b.

near-chance performance).² In addition to the test segment itself, a thin gray border was presented parallel to the segment (see Fig. 3). This border indicated the inside of the shape and served two purposes: it decreased the difficulty of the task by cueing observers to the general area on the probe shape on which to search for the test segment. It also permitted a correct figure/ground assignment, offsetting any bias observers may have to see the isolated segment as convex.

2.1. Experiment 1a: Negative minima vs. positive maxima

Experiment 1a compared the role of different types of curvature extrema in shape segmentation. Accuracy was measured for trials in which test segments were bounded on either side by positive maxima, or by negative minima, of curvature. If part decomposition is determined by negative minima as candidate boundaries, as postulated by the minima rule, accuracy for minima-bounded segments should be substantially greater than for maxima-bounded segments. Positive maxima constitute the best comparison case for testing the role of negative minima, since the two types of curvature extrema are locally equivalent, differing only in their sign of curvature. Moreover, by Attneave's analysis, they are postulated to have the same informational content.³ This experiment thus investigated whether curvature extrema play different roles in shape segmentation, based on their sign of curvature.

2.1.1. Methods

2.1.1.1. Observers. Twelve Rutgers University undergraduates participated in exchange for course credit. All were naïve to the purpose of the experiment.

² It should be noted that this presentation order only strengthens the logic on which the SVT rests, namely, that segments constituting natural units of shape for the visual system will produce better verification performance. In all trials, the observer has precise information about the segment before the shape appears. Greater failure in one condition to verify the segment's identity within the shape indicates a greater degree of incongruence between the perceptual parsing of the shape and the artificially divided segment.

³ Both extrema types are local peaks of curvature magnitude. However, for randomly-generated closed shapes, positive maxima tend on average to have greater curvature magnitude.

2.1.1.2. Stimuli. Probe shapes were random 2D smooth shapes presented as outlines. Shapes were generated online thus resulting in a unique set of shapes for each observer. This technique was employed to create a wide stimulus set. Each shape was first randomly generated as a 12-point polygon measuring between 0.86° and 4.28° of visual angle in diameter. The location of each point on the polygon was selected by projecting 12 radial axes, initially separated equally in polar angle (by 30°), from the center of the screen. In order to create variability in shapes, a random angular offset was applied to each vertex (uniformly distributed on $[-5^\circ, 5^\circ]$). This resulted in an angular separation ranging between 20° and 40° between neighboring points. The radial distance of each point was drawn from a uniform distribution between 0.43° and 2.14° of visual angle. A smoothing operation was then applied to each shape by fine-sampling the closed contour and convolving with a 1D Gaussian.

Two versions were created for each shape—the probe shape and the distorted shape. The distorted shape was generated by applying a displacement of set magnitude, but random direction, to each vertex in the initial polygon defining the probe shape. Four distortion magnitudes were used. Each displaced vertex was tested for a possible reversal in the sign of curvature (i.e., a switch from convex to concave, or vice versa) in the contour, resulting from the distortion procedure. Distortions resulting in such reversals were rejected; a new distorted shape was then sampled using the same procedure.

Curvature extrema were identified as zero crossings of the third derivative of the contour defining shape's outline. Segments were generated by selecting two consecutive extrema of the same sign (either both positive maxima or both negative minima) with at least one intervening extremum of the opposite sign. Thus, each minima-defined segment contained at least one positive maximum between its minima endpoints, and vice versa. If a generated shape did not contain any suitable segment, it was discarded and a new shape was sampled.

For match trials, a test segment of the appropriate type was randomly selected from the probe shape, and presented in conjunction with the probe shape. For mismatch trials, test segments were selected from the corresponding section on the distorted version of the shape, and presented in conjunction with the probe shape. Probe shapes and test segments were presented as white contours, ~ 1 arcmin thick, against a black background. Test segments were presented with a border cue that indicated the direction of the 'inside' of the probe shape. The border cue was drawn as a parallel gray segment, drawn by moving each point on the test segment in the locally normal direction, toward the inside of the probe shape from which the segment was selected. It was 4 arcmin thick.

2.1.1.3. Design. Three variables were manipulated: *segment type* (minima-bounded segments, i.e., segments whose end-

points were negative minima of curvature, or maxima-bounded segments), *trial type* (match or mismatch), and *distortion level* (10, 14, 18, or 22 arcmin) in mismatch trials. Trials were presented in blocks of 32. In total, there were 20 blocks for a total of 640 experimental trials, preceded by two practice blocks.

2.1.1.4. Procedures. Observers were shown the following sequence on each trial (see Fig. 2): (i) a test segment (350 ms), (ii) a mask (500 ms), and (iii) the entire probe shape (presented until response). The test segment matched exactly a portion of the probe shape on 50% of trials. The observer's task was to indicate (yes/no) by keyboard press, whether or not the test segment matched a portion of the probe shape. Incorrect responses produced a feedback tone.

2.1.2. Results and discussion

Results indicated much greater response accuracy for minima-defined segment trials (mean accuracy = 75.6%, $SE = 1.7$) than maxima-defined segments (mean accuracy = 63.7%, $SE = 2.6$). Converted to d' this difference translated into a nearly 100% increase in sensitivity. This difference in sensitivity persisted across all levels of distortion. The mean sensitivity for minima-defined segments was 1.49 ($SE = .14$) as compared to maxima-defined segments (mean = .76, $SE = .17$). An analysis of variance yielded significant main effects for segment type ($F(1, 11) = 32.4$, $p < .0001$) as well as distortion level ($F(3, 33) = 18.0$, $p < .0001$). Not surprisingly, discrimination performance improved with increasing distortion level. Moreover, there was a significant interaction ($F(3, 33) = 5.0$, $p < .01$). As is evident in Fig. 4, the difference in performance between positive maxima and negative minima is greater for larger distortions.

Observers were thus far more accurate to correctly verify the presence of contour segments taken from probe shapes if those segments were bounded by negative minima of curvature, than if they were bounded by positive maxima of curvature. Broadly speaking, these results suggest first that subsets of a shape are automatically segmented

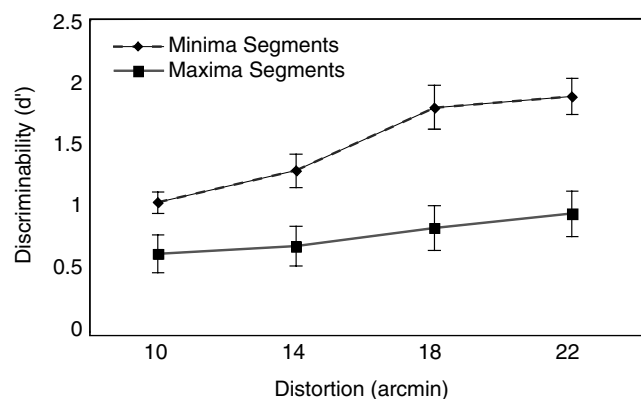


Fig. 4. Results of Experiment 1a. Error bars indicate standard error.

by human vision. Observers were not overtly asked to demarcate or otherwise indicate parts, but simply to perform the task of verifying whether or not a segment belongs to a shape. The fact that one type of contour segment—as determined by the geometric properties of its boundaries—was more verifiable than the other, suggests that one type of segment approximates more closely the natural units of shape as defined by the visual system. Specifically contour fragments with negative-minima endpoints were favored over those with positive-maxima endpoints.

2.2. Experiment 1b: Segment length

Shapes necessarily have more total positive curvature along their occluding contour than negative curvature. A contour with more total negative curvature than positive curvature would not be closed. One consequence of this geometric necessity is that segments of positive curvature on a closed shape are frequently longer than those of negative curvature. While the stimuli in Experiment 1a had the benefit of being generated entirely at random—thereby creating a wide class of shapes—a byproduct of such generative freedom was that minima-bounded segments were on average longer than maxima-bounded segments. This geometric inequality may lead to an alternative explanation of the data: perhaps the higher performance for minima-bounded segments simply reflects their increased lengths.

In order to investigate this possibility, Experiment 1b performed a controlled modification of the first experiment in which test segments in all conditions were constrained to have the same mean and range of lengths. This control was carried out by creating a balanced library of 640 pre-computed probe-shape and test-segment pairs.

2.2.1. Methods

Methods were identical to those of Experiment 1a, with the exception that trial stimuli (probe shapes and test segments) were pre-selected for segment length from shapes randomly generated using the same procedure as the first experiment. Specifically, stimuli were selected so that the set of minima-defined segments had the same mean as the set of maxima-defined segments ($=90$ arcmin), and a common range of ± 26 arcmin. Trials were presented in random order.

2.2.1.1. Observers. A new group of 12 Rutgers University undergraduates participated in exchange for course credit. All were naïve to the purpose of the experiment.

2.2.2. Results and discussion

The differential pattern of results in Experiment 1b was very similar to that of Experiment 1a, though there was a small overall decrease in accuracy across conditions. Mean accuracy for minima-defined segment trials was 69.9% ($SE = 2.4$) as compared to 58.5% ($SE = 2.4$) for maxima-




Fig. 5. Results of Experiment 1b. Error bars indicate standard error.

defined segments. Converted to d' , this translated to a mean sensitivity of 1.11 ($SE = .16$) for minima-defined segments and .47 ($SE = .16$) for maxima-defined segments. Once again there was a significant main effect of segment type ($F(1, 11) = 42.9, p < .0001$) as well as distortion level ($F(3, 33) = 12.9, p < .0001$). Their interaction was not statistically significant (Fig. 5).

It is clear from the results of Experiment 1b that the minima-segment advantage cannot be attributed simply to segment length. In both Experiments 1a and 1b, this advantage was robust, and obtained across all levels of distortion tested.

2.3. Experiment 2: Curvature extrema vs. positive curvature

As noted in the Introduction, single-cell studies in area V4 point to a *piecewise* coding of shape in the primate visual cortex, with different units within a population responding to either convex or concave segments of the occluding contour, at different locations on a shape (Pasupathy & Connor, 1999, 2001). One point of difference, however, between these findings and Hoffman & Richards' *minima rule* is that, in a coding scheme where individual units respond to purely convex or purely concave segments, the segmentation of the occluding contour would occur not at negative minima of curvature, but at inflection points—points where the occluding contour switches from being concave to convex, or vice versa (see Fig. 1). This reliance on inflection points is consistent with the idea that the dichotomy between positive and negative curvature is fundamental in shape representation: regions of positive Gaussian curvature on smooth 3D objects project to convex segments on the occluding contour, and tend to have a “thing-like” perceptual quality, whereas regions of negative Gaussian curvature project to concave segments on the occluding contour, and tend to have a “glue-like” quality (Koenderink, 1984; Koenderink & Van Doorn, 1982). These considerations suggest that inflection points should play a key role in the visual segmentation of occluding contours. Indeed, some models in the computer vision literature have explicitly proposed segmenting shapes based on the convexity of the candidate parts (Latecki & Lakamper, 1999; Rosin, 2000;

Vaina & Zlateva, 1990). The precise way in which convexity is used differs from model to model. Some have proposed that partitioning be executed in a way that results in parts having only convex regions (e.g., Latecki & Lakamper, 1999), whereas others propose overall maximization of part convexity for the entire shape decomposition scheme (e.g., Rosin, 2000).

Although contour segments bounded by negative minima often tend to have largely positive curvature, it is evident that for smooth shapes they must contain sections of negative curvature as well (since, by definition, negative minima themselves lie in negative-curvature segments of the occluding contour). Which factor determines visual part decomposition? Does the visual system explicitly use curvature extrema (specifically, negative minima) for segmentation, or does it partition shapes in a way that preserves uniformly positive-curvature regions? In order to differentiate between these two accounts, Experiment 2 compared segments bounded by extrema endpoints with segments defined by uniform sign of curvature, bounded by inflection endpoints. Since these inflection points always occurred inside pairs of consecutive extrema of the same sign of curvature, the resulting segments were always shorter than the extrema-endpoint segments. Therefore, it was necessary to include another inflection-endpoint condition in order to examine the influence of segment length. Thus we included segments that were bounded by inflections immediately outside of a pair of consecutive extrema. In this way, we were able to examine the effect of increasing length between these landmarks. In total there were six segment types (see Fig. 6), three defined with respect to minima-bounded segments and three defined with respect to maxima-bounded segments. For each, the three segments included the original extrema segments, as well as two types of inflection-bounded segments—one contained strictly within the extrema-defined segment (*inflection-shorter* segments) and one strictly containing the extrema-defined segment (*inflection-longer* segment).

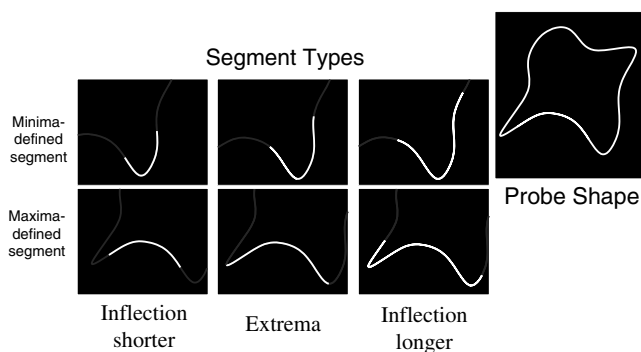


Fig. 6. Segment types for Experiment 2. Only the white contour was presented.

2.3.1. Methods

2.3.1.1. Observers. Fourteen Rutgers University undergraduates participated in exchange for course credit. All were naïve to the purpose of the experiment.

2.3.1.2. Stimuli. Probe shapes were generated using the same random-generation procedure as before. Test segments bounded by inflection points were generated by first selecting a segment bounded by two extrema of the same type (using the same procedure as before), and then either proceeding to the points of zero curvature inside the extrema endpoints (for *inflection-shorter* segments) or to the points of zero curvature outside both extrema (for *inflection-longer* segments). As in Experiment 1, probe shapes were generated by applying a smoothing operation to the bounding contours of randomly-generated polygons. The level of smoothing was much greater than in Experiment 1 in order to create meaningful inflection points. However, extended segments of zero curvature (rather than a single point) sometimes remained. In these cases, the inflection was taken to be the midpoint of the zero curvature segment.

Six types of segments were used: minima-bounded segments, inflection segments—shorter than minima, inflection segments—longer than minima, maxima-bounded segments, inflection segments—shorter than maxima, inflection segments—longer than maxima (see Fig. 6).

2.3.1.3. Design. Experiment 2 manipulated four variables: segment base type (minima-based vs. maxima-based), segment length (inflections-shorter, extrema, inflections-longer), trial type (match vs. mismatch), and distortion level (14, 18, 22 arcmin). Twenty blocks of 36 trials were presented to each observer, preceded by two practice blocks.

2.3.1.4. Procedure. The stimulus presentation and trial sequence were identical to Experiments 1a and 1b.

2.3.2. Results

Fig. 7 shows performance for all segment types and distortion levels. An ANOVA revealed a significant main effect for *segment base type* (minima-based or maxima-based) ($F(1,13) = 21.64, p < .0001$), namely, a superior performance for minima-based segments over their maxima-based counterparts. There was a main effect of *segment length* ($F(2,26) = 10.24, p < .001$) showing an overall increase in performance with increasing segment length. Finally, there was a significant effect of *distortion level* ($F(2,26) = 36.89, p < .0001$). The interaction between segment base type and segment length was also significant ($F(4,44) = 5.48, p = .01$).

Since the influence of distortion level was consistent across all six types of segments, we focus on the interaction between segment base type and segment length. Fig. 8




Fig. 7. Results of Experiment 2. Data from all six length conditions across three distortion levels. Error bars indicate standard error. A clear advantage in d' was observed for the minima-bounded segments and the longer-than-minima inflection segments across all distortion levels as compared to all four other length conditions.

shows the data collapsed across distortion levels. Performance for all three minima-based segments was superior than that for all maxima-based segments. The highest performance was observed for the minima-bounded segments and the longer-than-minima inflection segments.

Focusing on the three minima-based segments, we find an asymmetric differential in performance, in going from the minima-bounded segments to the inflections-shorter vs. inflection-longer segments. There is a large drop in performance in going from the minima-bounded segments to the inflections-shorter segments. However, the performance for the inflections-longer segments is no better than that for the minima-bounded segments. This asymmetry indicates that differences in accuracy were not due to differences in segment length per se, but rather due to the presence of negative minima. In other words, performance increased to the extent that a minima-defined unit was present. Furthermore, the leveling off is unlikely to be due to a ceiling effect. All three distortion levels (collapsed in Fig. 8)

exhibit the same pattern of leveling off from minima to longer than minima inflections, even though they have different accuracy levels. Finally, the raw accuracy data show that no condition in Experiment 2 (for either match or mismatch conditions) yielded a mean accuracy higher than 82%, thus making a ceiling effect unlikely.

It is also noteworthy that the asymmetry found for minima-based segments was not observed for the maxima-based segments; this curve shows a roughly linear (though barely significant) increase in accuracy with increasing segment length. Fig. 9 shows these differentials in performance (“slopes”)—from inflection-shorter to extrema segments, and extrema segments to inflection-longer segments—for both minima-based and maxima-based segments. This plot clarifies that a significant difference in slope is observed only for the minima-based segments. These results thus provide further evidence that

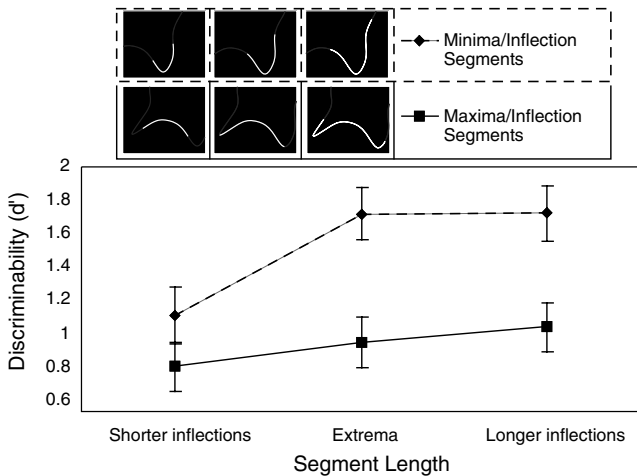


Fig. 8. Results of Experiment 2, collapsed over distortion level, showing the interaction between segment base type and segment length.




Fig. 9. Illustration of the differentials in performance (“slopes”)—from inflection-shorter to extrema segments, and extrema segments to inflection-longer segments—for both minima-based and maxima-based segments. Lines represent the mean slope for each extrema type.

negative minima provide the strongest contour-based cues to shape segmentation.

2.3.3. Discussion

The results of Experiment 2 indicate that segments with negative-minima boundaries represent the most identifiable segment units. Simply put, performance increased to the extent that a minima-defined unit was present in the test segment. Performance increased with segment length from shorter inflection segments to minima-bounded segments. However, performance for the minima and longer-than-minima inflection segments was virtually identical—consistent with the fact that these two segment types contain the same ‘natural part’ unit.

The current results show no evidence of an explicit role of inflections in segmentation. Segments of uniform positive curvature were not as readily identified as minima-defined segments. Uniform positive-curvature or strict convexity thus does not appear to be the key geometric factor in part segmentation. Rather, negative minima of curvature determine decomposition.⁴

3. Graded segmentation and part salience

Although the preceding evidence has brought us closer to understanding the way in which part boundaries are visually defined, our understanding of the role of parts within a global representation of complex shape remains limited. Complex objects may contain many parts. Within a shape’s global representation, the status of a given shape part is clearly dependent upon many factors including scale, object complexity, and other local and global geometric determinants. Visual experience suggests a continuum of part importance, upon which mere texture-like surface variations are given low emphasis, and prominent parts are awarded high emphasis. It has long been suggested that complex shape representations are organized as hierarchical tree structures (Marr & Nishihara, 1978; Palmer, 1977; Rom & Medioni, 1993). More stable or more salient parts are represented at higher levels (near the root of a hierarchy tree), while smaller or less stable parts are represented at lower-levels. However, hierarchical structure provides only the broadest framework for understanding object representation: Even parts at the same hierarchical level, and of the same size, can differ a great deal in their perceptual salience.

⁴ Although the results of Experiment 2, provide evidence against part segmentation based on uniformly positive curvature (i.e., convex-only segments), or parsing at inflection points, they do not automatically rule out models based on the maximization of convexity for the entire part decomposition scheme (e.g., Rosin, 2000). Such decomposition methods also often result in part boundaries located at or near negative minima of curvature (although they tend to perform poorly on shapes with curved axes—because such shapes receive low convexity measures, despite being perceived as a single part).

Hoffman and Singh (1997) proposed that the representation of a part is graded, and that the visual *salience* of a part is influenced by at least three geometric factors: its size relative to the whole object, its degree of protrusion (as defined by the ratio of its perimeter to the width of the part at its base), and the strength of its boundaries (as characterized by both the magnitude of curvature at the part boundaries, and the turning angle between the two nearest inflections on either side of a boundary point). This gradedness of part representations is also emphasized by a computational theory of shape representation by Siddiqi, Kimia, Tannenbaum, and Zucker (2001). Siddiqi et. al proposed that shapes can be described via a combination of three continuous processes—bending, protruding, and partitioning. The application of these processes generates a 3D space within which a shape’s description may be localized. Though both these schemes have strong intuitive appeal, little direct psychophysical support exists for the way in which individual geometric factors influence the strength or salience of a part’s representation (but see Hoffman & Singh, 1997 and Siddiqi et al., 2001, for influences mediated by figure-ground perception and visual search, respectively). The following experiments directly measure the contributions of boundary strength, contour length, base-width, and part protrusion to a part’s visual salience.

3.1. Segment identification task

The yes/no segment-verification method used in the previous experiments required a distortion in the shape of the test segment (for the mismatch or “no” trials). Unlike Experiments 1 and 2, which compared performance across segments with qualitatively different boundary cues, Experiments 3 and 4 manipulated quantitative properties of test segments—such as the turning angles at their part boundaries. The random distortion method used in SVT does not permit tight control over these properties across both versions of the segment (the original and the distorted). It was therefore necessary to modify the experimental method to a forced-choice task that did not involve segment distortion.

The 4AFC segment-identification task (SIT) employed in Experiments 3 and 4 allows direct measurement of the influence of individual geometric factors on the representational strength of a given part. Observers were shown an outline shape followed by a contour segment taken from that shape, and asked to identify the position of the segment within the shape by choosing between four screen quadrants (see Fig. 10). The experimental rationale for the SIT was that identification performance should be superior for contour segments that are more strongly (or more likely to be) emphasized in a shape’s representation. By systematically manipulating various geometric properties of parts, we can better understand the role of parts in shape representation, and their geometric determinants.




Fig. 10. The trial sequence for the segment-identification task used in Experiments 3 and 4.

3.2. Experiment 3: Boundary strength

In Experiments 1 and 2, we demonstrated that, in the context of closed contours bounding 2D shapes, the sign of curvature plays an important role in shape segmentation. In particular, negative minima are more effective in creating boundaries between parts than positive maxima. In the current experiment, we examined the role of curvature magnitude at those negative minima, namely, their ‘boundary strength’ (Hoffman & Singh, 1997). We investigated whether parts become more identifiable when their boundaries are characterized by high negative-curvature as opposed to low-negative curvature. Observers performed the segment-identification task for a set of shapes and test segments that varied in the turning angle at their boundaries, and their contour length (see Fig. 11).

3.2.1. Methods

3.2.1.1. Observers. Sixteen Rutgers University undergraduates participated in exchange for course credit. All were naïve to the purpose of the experiment.

3.2.1.2. Stimuli. Study shapes in Experiments 3 and 4 were first created as randomly-generated polygons similar to those in Experiments 1 and 2. However, a lower degree of smoothing was applied in order to permit tighter control over the turning angles. In order to generate test segments with specific values of turning angle at the boundaries, two steps were employed. For each shape, a test segment, defined by the contour falling between two consecutive negative-minima, was selected at random. The shape was then modified locally to set the turning angles at the boundaries of the segment to the required value. Specifically, each endpoint was translated along the bisector of the angle formed by its neighbors on either side, so that the angle formed by the three points had the required value. The resulting polygonal shape was then smoothed using a 1D Gaussian convolution as in Experiment 1. This technique was employed to produce a large library of randomly-generated shapes for each turning-angle condition. After shape generation, test segments were measured for length. Based on these two properties, 500 shape–segment stimulus pairs were selected such that there were 50 shapes for each of the 5×2 conditions.

When test segments were presented in isolation, they were rotated so that they always pointed upward, i.e., with their midpoint extending into the upper middle screen quadrant. Each segment was presented just above the center of the screen. This technique prevented the segment’s orientation or position from providing a cue as to its original location on the shape.




Fig. 11. Typical stimuli from Experiment 3. The ten shapes represent the ten experimental conditions. Dark areas reflect the test segments. In the experiment, the shapes and segments were presented as white contours against a black background. Lengths measure the unfolded contour segment.

3.2.1.3. Design and procedure. Observers were presented the following sequence on each trial: (1) the study shape for 1000 ms, (2) a mask 500 ms, (3) test segment for 500 ms, (4) a mask 500 ms, and (5) a screen divided into four quadrants presented until response (see Fig. 10). Their task was to select, by key press, the quadrant within which the segment appeared on the probe shape. Incorrect responses received auditory feedback.

Two independent variables were manipulated: segment length (with mean and ranges of: 137 ± 11 arcmin and 68.5 ± 11 arcmin), and the turning angle at the segment's boundary points (15° , 30° , 45° , 60° , and 75°). Thus, ten types of segment were possible (see Fig. 11). Fifty unique shapes were generated for each of the 10 segment types, resulting in 500 trials. These trials were presented in 20 blocks of 25 each, and preceded by 30 practice trials.

3.2.2. Results and discussion

Significant effects were obtained both for segment length ($F(1,15) = 61.25$, $p < .0001$) and for turning angle at part boundaries ($F(4,60) = 33.44$, $p < .0001$), with a marginal interaction between the two ($F(4,60) = 2.50$, $p = .05$). The results thus demonstrate a systematic influence of boundary angle and segment length on segment identification (see Fig. 12). A monotonic increase in accuracy with boundary strength is evident in both segment length conditions. As test segments were characterized by increasing turning angles at their boundaries, observers were more accurate to identify them. A roughly 12% increase in performance was observed for longer segments as compared to shorter segments, across boundary angles.

3.3. Experiment 4: Segment length and protrusion

In addition to the influence of turning angle at part boundaries, the results of Experiment 3 exhibited a systematic increase in identification performance with increasing segment length. This is not surprising: the longer the test segment, the greater its percentage of the occluding contour of the study shape. Any contour-based shape representation scheme would thus predict a systematic increase in performance with increasing segment length. However, a description of just the occluding contour is unlikely to serve as the final representational format for shape. A great deal of evidence suggests that the visual system represents not just the contour (i.e., a 1D “string” representation), but the region bounded by the contour, e.g., using a skele-

ton or axis-based description (Blum, 1973; Blum & Nagel, 1978; Brady & Asada, 1984; Burbeck & Pizer, 1995; Feldman & Singh, 2006; Fulvio & Singh, 2006; Kovacs, Feher, & Julesz, 1998; Leyton, 1988; Marr, 1977; Marr & Nishiura, 1978; Sebastian & Kimia, 2005; Singh & Hoffman, 2001).

In Experiment 4, we investigated whether a region-based property—namely, part protrusion—exerts a systematic influence on segment identification performance. Defined as the ratio of a part's perimeter (or segment length) to the length of its “base” (the straight-line join of its two boundary points), protrusion provides a measure of the extent to which a part ‘sticks out’ from a shape (Hoffman & Singh, 1997). Parts with equal contour length may vary widely in their protrusion. In Experiment 4, we examined whether part protrusion provides a better predictor of segment identification performance than sheer contour length. This was done by manipulating two parameters of the test segment: contour length and base width (measured as the distance between the two boundary points). For a fixed segment length, increasing its base width decreases the part's protrusion. For a given base width, increasing segment length increases the part's protrusion. The factorial design obtained by crossing contour length with base width thus allows an evaluation of the individual influences of both variables—and hence the influence of part protrusion, which depends on both.

3.3.1. Methods

3.3.1.1. Observers. Twelve Rutgers University undergraduates participated in exchange for course credit. All were naïve to the purpose of the experiment.

3.3.1.2. Stimuli. Study shapes and test segments were generated using a similar technique as in Experiment 3. Base lengths were manipulated by translating the endpoints of the test segment on the study shape equally, either toward or away from each other, along their straight-line join. Segment lengths were selected in the same manner as described in Experiment 1b. Although the turning angle at boundary points was not a manipulated variable, turning angles were constrained to lie between 45° and 65° .

3.3.1.3. Design and procedure. The trial sequence and the task of the observer were identical to Experiment 3. Two independent variables were manipulated: segment length (three ranges: 80 ± 11 arcmin, 125 ± 11 arcmin, and 170 ± 11 arcmin), and base width (24, 37, 50, and 62 arcmin). A total of 12 types of test segments were thus used (see Fig. 13). Fifty unique shapes were generated for each segment type, resulting in 600 experimental trials. These trials were presented in 20 blocks of 30 each, preceded by 30 practice trials.

3.3.2. Results and discussion

As expected, segment length exerted a significant influence on identification performance ($F(2,22) = 89.28$,




Fig. 12. Accuracy results from Experiment 3. Error bars indicate standard error.




Fig. 13. Representative stimuli from Experiment 4, one for each combination of segment length and boundary turning angle. The darkened contour in each depicts a test segment.

$p < .0001$), with longer segments yielding better performance. In addition, the part's base width also exhibited a significant main effect ($F(3,33) = 91.56$, $p < .0001$), with longer base widths yielding *poorer* identification performance. The interaction between segment length and base width was also significant ($F(6,66) = 5.89$, $p < .0001$). As noted above, this is the pattern of performance one would expect based on the influence of part protrusion.

A simple regression revealed that segment length accounted for 40.6% of the response variance. However, when base width was added to the regression model, even in linear form, the combination of segment length and base width accounted for 65.9% of the variance—a statistically significant boost in R^2 ($p < .0001$).

One may generally expect wider bases to fall on larger parts. The decrease in performance associated with increase in base width indicates, however, that the increase in performance was not simply due to a influence of part size, but rather due to an influence of increasing part protrusion, as a result of decreasing base width (Fig. 14).

At a first level of analysis, the results reveal three influential factors in determining part salience: contour length, base width, and part protrusion. The results also suggest, however, that while contour length is a natural predictor of part identification, part protrusion (based on a combination of contour length and base-width) provides a stronger predictor. The finding that decreasing base width increases part identification suggests that parts are not more salient simply based on the percentage of the stimulus they comprise. Rather, properties of the part's region, such as its protrusion, are also strongly determinant in a part's representational strength.




Fig. 14. Accuracy results for Experiment 4. Error bars indicate standard error.

4. General discussion

We have presented a new method for studying geometric influences on part-based representation of shape. By employing novel and randomly-generated shapes in a performance-based task, we quantified the influence of shape geometry on part segmentation in the absence of naming or recognition of familiar shapes. Segment identification was found to depend on qualitative geometric properties of the segments' endpoints, as well as on metric properties of the bounded region. Negative minima provided the best contour-based cue to segmentation: segments with negative-minima endpoints were identified with greater accuracy than those with comparable positive-maxima endpoints. When segments with curvature-extrema endpoints were compared to those with inflection endpoints, segments with negative-minima endpoints were again found to elicit superior identification. This superiority was not an artifact of increased segment length: increasing segment length beyond the negative-minima endpoints of a segment did not result in additional increase in performance. These findings also suggest that uniform sign of curvature does not provide the best characterization of perceptually natural parts. Finally, the strength of part representation was shown to be graded: increasing segment length, turning angles at part boundaries, and part protrusion all resulted in a systematic increase in segment identification performance.

4.1. Structural shape representation

There is general consensus that the visual system segments object shapes into component structural units and organizes the representation of shape in terms of these units. However, which structural representation scheme is employed by vision remains a topic of some debate. Given evidence from neurophysiology (Pasupathy & Connor, 1999, 2001), it seems possible that shape is represented by encoding curvature properties of its bounding contour.

As discussed earlier, however, many have also argued for a more region-based representation scheme that incorporates, in addition to contour properties, those of the enclosed region as well—perhaps by employing an axis or skeleton-based representation (Blum, 1973; Feldman & Singh, 2006; Sebastian & Kimia, 2005). Such a scheme allows one to capture additional shape properties such as local and global symmetries.

The differential performance between positive- and negative-curvature segments throughout these experiments itself suggests that representation of shape was not simply contour-based. Segments in these two conditions were largely identical in terms of their local properties. In Experiment 2, for instance, the influence of added contour length was shown to be dependent on the type of surrounding segment. The fact that observers exhibit differential performance with these segments suggests that representation of the contour is necessarily tied to the region it encompasses. In this sense, it is perhaps not entirely apt to treat the sign of curvature as a property simply of a contour—since it necessarily alters region-based descriptors of shape. Furthermore, the results of Experiment 4 demonstrate that observers are sensitive to region-based properties of shape not directly associated with the 1D representation of contour, such as a part's base width (or cut length) and its degree of protrusion. Clearly, the differential influence of such metric properties is not simply reducible to decreased grouping or continuation as in studies of contour integration. Such findings, though, are only the beginning of an understanding of region-based representation.

4.2. Part independence

In Cohen and Singh (2006), we showed that geometric properties determining part salience also determine the extent to which a given part can be integrated into a *global* estimate of a shape. For instance, when a part was characterized by shallow boundaries, it strongly influenced judgments of the shape's global orientation. However, as a part's boundaries became sharper, and the part more distinct, its influence became less integrated into the global estimate of the shape's orientation. These results suggested that the continuum of part salience may be naturally characterized as one of *part independence*. The extent to which a part constitutes a salient sub-unit also determines the extent to which it is incorporated into a unified global representation of the shape—and hence the extent to which its influence is separable from the rest of the shape in making a global shape estimate (such as location or orientation).

4.3. Salience and spatial scale

One explanation for the increased part salience due to contour length, boundary strength, and part protrusion relates to responses across multiple spatial scales (see, e.g., Burbeck & Pizer, 1995, for the influence of scale on

shape representation). Specifically, it is possible that a part's perceptual salience may be characterized in terms of the range of spatial scales at which the part can be detected. For instance, very shallow turning angles at negative minima will only be detectable—i.e., coded as distinct from a straight contour—at a fine spatial scale. As the turning angle gets sharper, however, even coarse scale detectors will respond to the curvature extrema—whether this is construed in terms of the responses of end-stopped cells (see Dobbins, Zucker, & Cynader, 1987, 1989), or simply due to a weakening in contour integration due to the large turning angles between neighboring edges along the contour. A part with high protrusion will similarly tend to be detectable at a wider range of spatial scales (specifically, even at coarser scales) than one with low protrusion.

Is spatial scale alone sufficient to explain the full spectrum of geometric influences on part salience? There seem to be some hierarchical aspects of part perception that are not easily reducible to spatial scale. For instance, a part may be more salient, or receive greater emphasis in shape representation, if it is perceived as a 'base' part (such as a 'trunk') from which other parts (such as 'branches') protrude—even if that 'base' part is smaller in size. Furthermore, parts may vary in salience depending on the number of surrounding similar parts (the tines on a fork vs. the teeth of a comb). There is no doubt that the multi-scale nature of basic visual processing can account for many aspects of shape perception, and of properties that determine a part's perceptual status within a shape. However, a full understanding of part-based representation is likely to involve other, more structural and hierarchical, characteristics as well. The experimental techniques developed here provide a tool that can provide insight into the geometric influences on part segmentation and part salience. Uncovering how these geometric properties are encoded by neural mechanisms remains a challenge for future work.

Acknowledgments

We thank Doug DeCarlo, Jacob Feldman, Eileen Kowler, and Johan Wagemans for helpful comments and suggestions. This work was funded by NSF BCS-0216944.

References

- Attneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, *61*, 183–193.
- 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). One-sided edge assignment in vision. 1. Figure-ground segmentation and attention to objects. *Current Directions in Psychological Science*, *4*, 140–146.

- 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.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115–147.
- Biederman, I., & Cooper, E. E. (1991). Priming contour-deleted images: Evidence for intermediate representations in visual object recognition. *Cognitive Psychology*, *23*, 393–419.
- Blum, H. (1973). Biological shape and visual science (Part I). *Journal of Theoretical Biology*, *38*, 205–287.
- Blum, H., & Nagel, R. N. (1978). Shape description using weighted symmetric axis features. *Pattern Recognition*, *10*, 167–180.
- Brady, M., & Asada, H. (1984). Smoothed local symmetries and their implementation. *The International Journal of Robotics Research*, *3*, 36–61.
- 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*(4), 313–321.
- Cohen, E. H., & Singh, M. (2006). Perceived orientation of complex shape reflects graded part decomposition. *Journal of Vision*, *6*(8), 805–821.
- Denisova, K., Singh, M., & Kowler, E. (2006). The role of part structure in the perceptual localization of a shape. *Perception*, *35*, 1073–1087.
- Desimone, R., Albright, T., Gross, C. G., & Bruce, C. (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. *Journal of Neurophysiology*, *4*, 2051–2062.
- 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*, *99*(3), 275–325.
- Dobbins, A., Zucker, S. W., & Cynader, M. S. (1987). Endstopped neurons in the visual cortex as a substrate for calculating curvature. *Nature*, *329*(6138), 438–441.
- Dobbins, A., Zucker, S. W., & Cynader, M. S. (1989). Endstopping and curvature. *Vision Research*, *10*, 1371–1387.
- Feldman, J. (1997). Curvilinearity, covariance, and regularity in perceptual groups. *Vision Research*, *37*(20), 2835–2848.
- Feldman, J., & Singh, M. (2005). Information along contours and object boundaries. *Psychological Review*, *112*(1), 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.
- 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.
- Fulvio, J., & Singh, M. (2006). Surface geometry influences the shape of illusory contours. *Acta Psychologica*, *123*, 20–40.
- 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.
- Gross, C., Rocha-Miranda, C. E., & Bender, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the macaque. *Journal of Neurophysiology*, *35*, 96–111.
- Hayward, W., Tarr, M., & Corderoy, A. (1999). Recognizing silhouettes and shaded images across depth rotation. *Perception*, *28*, 1197–1215.
- Hoffman, D. D., & Richards, W. A. (1984). Parts of recognition. *Cognition*, *18*, 65–96.
- Hoffman, D. D., & Singh, M. (1997). Saliency of visual parts. *Cognition*, *63*, 29–78.
- Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurons in the cats striate cortex. *Journal of Neurophysiology (London)*, *148*, 574–591.
- Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cats visual cortex. *Journal of Neurophysiology (London)*, *160*, 106–154.
- Hulleman, J., te Winkel, W., & Boselie, F. (2000). Concavities as basic features in visual search: Evidence from search asymmetries. *Perception & Psychophysics*, *62*, 162–174.
- Koenderink, J. (1984). What the occluding contour tell us about solid shape. *Perception*, *13*, 321–330.
- Koenderink, J., & Van Doorn, A. (1982). The shape of smooth objects and the way contours end. *Perception*, *11*, 129–137.
- 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.
- 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.
- Latecki, L., & Lakamper, R. (1999). Convexity rule for shape decomposition based on discrete contour evolution. *Computer Vision and Image Understanding*, *73*, 441–454.
- Leek, E. C., Reppas, I., & Arguin, M. (2005). The structure of three-dimensional object representations in human vision: Evidence from whole-part matching. *Journal of Experimental Psychology: Human Perception and Performance*, *31*, 668–684.
- Leyton, M. (1988). A process grammar for shape. *Artificial Intelligence*, *34*, 213–247.
- Leyton, M. (1989). Inferring causal history from shape. *Cognitive Science*, *13*, 357–387.
- Liu, Z., Jacobs, D., & Basri, R. (1999). The role of convexity in perceptual completion: Beyond good continuation. *Vision Research*, *39*, 4244–4257.
- Marr, D. (1977). Analysis of occluding contour. *Proceedings of the Royal Society of London. Series B*, *197*, 441–475.
- Marr, D., & Nishihara, H. K. (1978). Representation and recognition of three-dimensional shapes. *Proceedings of the Royal Society of London. Series B*, *200*, 269–294.
- 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. (1977). Hierarchical structure in perceptual representation. *Cognitive Psychology*, *9*, 441–474.
- Pasupathy, A., & Connor, C. E. (1999). Responses to contour features in macaque area V4. *Journal of Neurophysiology*, *82*, 2490–2502.
- 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 area V4. *Nature Neuroscience*, *5*(12), 1332–1338.
- Perrett, D. I., Rolls, E. T., & Caan, W. (1982). Visual neurones responsive to faces in the monkey temporal cortex. *Experimental Brain Research*, *47*, 329–342.
- Pettet, M. W., McKee, S. P., & Grzywacz, N. M. (1998). Constraints on long range interactions mediating contour detection. *Vision Research*, *38*(6), 865–879.
- Phillips, F., Todd, J., Koenderink, J., & Kappers, A. (2003). Perceptual representation of visible surfaces. *Perception & Psychophysics*, *65*, 747–762.
- Richards, W. A., & Hoffman, D. D. (1985). Codon constraints on closed 2d shapes. *Computer Vision, Graphics, and Image Processing*, *31*, 156–177.
- Rom, H., & Medioni, G. (1993). Hierarchical decomposition and axial shape description. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *11*, 823–839.
- Rosin, P. (2000). Shape partitioning by convexity. *IEEE Transactions on Systems, Man, and Cybernetics—Part A. Systems and humans*, *30*(2), 202–210.
- Rubin, E. (1921). *Visuell wahrgenommene figuren*. Kobenhaven: Glydenalske boghandel.

- Sebastian, T. B., & Kimia, B. B. (2005). Curves vs skeletons in object recognition. *Signal Processing*, *85*, 247–263.
- 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., Tresness, K. J., & Kimia, B. B. (1996). Parts of visual form: Psychophysical aspects. *Perception*, *25*, 399–424.
- Singh, M., & Hoffman, D. D. (1998). Part boundaries alter the perception of transparency. *Psychological Science*, *9*, 370–378.
- 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, NY: Elsevier Sciences.
- Tanaka, K., Saito, H., Fukada, Y., & Moriya, M. (1991). Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *Journal of Neurophysiology*, *66*, 170–189.
- Tversky, B., & Hemenway, K. (1984). Objects, parts, and categories. *Journal of Experimental Psychology: General*, *113*, 169–193.
- Vaina, L., & Zlateva, S. (1990). The largest convex patches: A boundary-based method for obtaining object parts. *Biological Cybernetics*, *62*, 225–236.
- Vecera, S. P., Behrmann, M., & Filapek, J. C. (2001). Attending to the parts of a single object: Part-based selection of limitations. *Perception & Psychophysics*, *63*(2), 308–321.
- Waeytens, K., Hanouille, I., Wagemans, J., & d'Ydewalle, G. (1994). Human representation of closed contours. In F. Dillen, I. Van de Woestijne, & L. Verstraelen (Eds.), *Geometry and topology of submanifolds vi* (pp. 298–306). Singapore: World Scientific.
- Watson, S. E., & Kramer, A. F. (1999). Object-based visual selective attention and perceptual organization. *Perception & 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.