Visual representation of contour and shape

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

To appear in:
Oxford Handbook of Perceptual Organization
Oxford University Press
Edited by Johan Wagemans

Abstract

Contours provide an essential source of information about shape and, along contours, points with the greatest magnitude of curvature tend to be most informative. This concentration of information is closely tied to internal generative models of contours employed by the visual system. In going from open to closed contours, the sign of curvature becomes perceptually significant, with negative-curvature (concave) sections of a contour being more informative, and playing an important role in part segmentation. The visual system represents complex shapes by segmenting them into simpler parts ("simpler" because they have less negative curvature). Points of negative minima of curvature provide an important cue for part segmentation; however being entirely local and contour-based features, they cannot fully predict part segmentation. The visual system employs not only a contour-based representation of shape, but also a region-based one, making explicit properties such as axis curvature, local width of the shape, and locally-parallel and locally-symmetric structure. A region-based representation of shape based on Bayesian estimation of the shape skeleton provides a successful account of part segmentation. Moreover, psychophysical results from a variety of domains provide evidence for the representation of region-based geometry by human vision, based on the shape skeleton. Even at the level of so-called "illusory contours," nonlocal region-based geometry exerts a strong influence. We conclude that, as far as the visual representation of shape is concerned, contour geometry cannot ultimately be studied in isolation, but must be considered conjointly with region-based geometry.

Keywords: contours, information, internal models, generative models, curvature, parts, shape skeleton, axis, region-based geometry, Bayesian inference, MAP skeleton.

1. Contours and information

Images are far from uniform in their information content. Rather information tends of be concentrated in regions around contours. This makes good sense: the presence of a contour signals some physically significant "event" in the world—whether it be the occluding boundary of an object, a reflectance change, or something else. Indeed, human observers are just as good at scene recognition with line drawings as they are with full-color photographs (e.g., Walther et al., 2011). Similarly, object recognition (e.g., Biederman & Ju, 1988) and 3D shape perception (e.g., Cole et al., 2009) are often just as good with line drawings as they are with shaded images. It is therefore not surprising that line drawings have a long history—having been used by humans as an effective mode of visual depiction and communication since prehistoric times (as evidenced, for example, by the Chauvet cave paintings; see, e.g., Clottes, 2003).

In his seminal article, Attneave (1954) noted not only the high-information content of contours in images, but argued also that, along contours, points of maximal curvature carry the greatest information. In support of this latter claim, Attneave provided two lines of evidence. First, he briefly reported the results of a study in which participants were asked to approximate a shape as closely as possible with only a limited number of points, and then to indicate the locations corresponding to those points on the original shape. Histograms of locations selected by the participants exhibited sharp peaks at local maxima of curvature—pointing to their importance in shape representation. Second, Attneave made a line drawing of a sleeping cat using only local curvature maxima that were then connected with straight-line segments. The resulting drawing was readily recognizable as a cat (now famously known as "Attneave's cat"), suggesting that not much information had been lost.

Attneave's second line of evidence has been the subject of further discussion and some controversy; the precise result appears to depend on the geometry of the contour (whether or not it has large variations in curvature and salient maxima) and the presence of other types of competing candidates points (e.g., Kennedy & Domander, 1985; De Winter & Wagemans, 2008a; 2008b; Panis, De Winter, Vandekerckhove & Wagemans, 2008). His first, experimental, finding has been uncontroversial, however. Indeed, Norman et al. (2001) conducted a study along the lines described briefly in Attneave (1954) using silhouettes cast by natural 3D objects (sweet potatoes), and replicated his findings (see Figure 1a for sample results).[1] Similarly, De Winter & Wagemans (2008b) found that when participants are asked simply to mark "salient points" along the bounding contours of 2D shapes—without being required to replicate the shape—they are again most likely to pick local maxima of curvature. As we will see, curvature extrema play an important role in modern theories of shape representation as well (Hoffman & Richards, 1984; Richards, Dawson & Whittington, 1986; Leyton, 1989; Hoffman & Singh, 1997; Singh & Hoffman, 2001; De Winter & Wagemans, 2006; 2008a; Cohen & Singh, 2007).

But why should curvature maxima be the most informative points along a contour? The link between contour curvature and information content follows fairly directly from Shannon's theory of information (in particular, from the definition of surprisal as $u = -\log(p)$), once one adopts a simple and empirically motivated generative model of contours (Feldman & Singh,

[1] A detailed report of Attneave's original experiment was apparently never published. His own article cites only a "mimeographed note."

2005; Singh & Feldman, 2012).[2] Specifically, one may ask, as one moves along a contour, where is the contour likely to go "next" at any given point? A great deal of psychophysical work on contour integration and contour detection has shown that the visual system implicitly expects that a contour is most likely to go "straight" (i.e., to continue along its current tangent direction), and that the probability of "turning" away from the current tangent direction decreases monotonically with the magnitude of the turning angle (Field et al., 1993; Feldman, 1997; Geisler et al., 2001; Elder & Goldberg, 2002; Yuille, Fang, Schrater & Kersten, 2004). The visual system's local probabilistic expectations about contours may thus be summarized as a von Mises (or circular normal) distribution on turning angles, centered on 0 (see Figure 1b; Feldman & Singh 2005; Singh & Feldman 2012). Indeed, even the assumption of a specific distributional form is not necessary to derive Attneave's claim; all that is needed is that the distribution on turning angles peak at 0 degrees, and decrease monotonically on both sides. It then follows directly from this that the surprisal, $u = -\log(p)$, increases monotonically with the magnitude of the turning angle. And turning angle, of course, is simply the discrete analog of curvature. Hence maxima of curvature are also maxima of contour information—which is precisely Attneave's claim.

(a)

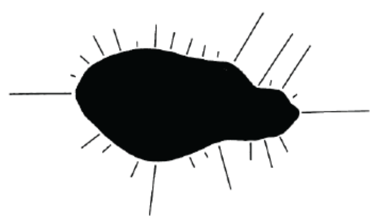

(b)

Figure 1. (a) Generative model of open contours expressed a probability distribution on turning angle from the current contour orientation. The distribution is centered on 0, meaning that going "straight" (i.e., zero turning) is most likely, with the probability decreasing monotonically with turning angle in either direction. This empirically motivated generative model explains why information along contour increases monotonically with curvature (Feldman & Singh, 2005). (b) Sample results from Norman et al.'s (2001) replication of Attneave's experiment. Histograms of points selected by subjects show peaks at maxima of curvature.

[2] Note that the formula for the surprisal is consistent with the simple everyday intuition that improbable events, when they occur, are cause for greater surprise—and hence are more informative—than when a highly probable, or expected, event occurs. As they say, "man bites dog" is news; "dog bites man" is not.

One can go further, however. Attneave (1954) treated curvature only as an unsigned quantity, i.e., simply as a magnitude. For a closed contour, however (such as the outline of an object), it is not only meaningful but also more appropriate to treat curvature as a signed quantity—specifically, as having positive sign in convex sections of the contour, and negative sign in concave sections. Indeed, there are principled reasons to expect that the visual system should treat convex and concave portions of a shape quite differently (Koenderink & van Doorn, 1982; Koenderink 1984; Hoffman & Richards, 1984). From the point of view of information content of contours, however, the key observation is that on closed contours, the probability distribution on turning angles is not centered on 0, but rather is biased such that positive turning angles (involving turns toward the shape, or figural side of the contour) are more likely than negative turning angles. Indeed, this must be the case if the contour is to eventually close in on itself. And it entails, via the -log(p) relation, an asymmetry in surprisal, such that negative curvature is more "surprising"—and hence more informative—than corresponding magnitudes of positive curvature (see Feldman & Singh, 2005, for details). This asymmetry in information content is supported by empirical findings showing that changes at concavities are easier to detect visually than corresponding changes at convexities (Barenholtz, Cohen, Feldman & Singh, 2003; Cohen, Barenholtz, Singh & Feldman, 2005), although there are nonlocal influences as well—based on, for example, whether a shape change alters qualitative part structure (e.g., Bertamini & Farrant, 2005; Vandekerckhove, Panis & Wagemans, 2008). (See also Section 4 for more on nonlocal influences in shape perception.)

In summary, Attneave's claim about curvature and information follows from a simple and empirically-motivated generative model of contours. And, as noted above, Attneave's theoretical claim can also be extended to closed contours, with the result that negative curvature segments carry more information that corresponding positive curvature segments.[3] The stochastic generative model of contours may also be extended to incorporate the role of co-circularity, i.e., the visual expectation that that contours tend to maintain their curvature (Singh & Feldman, 2012). Psychophysical evidence for this expectation by the visual system comes from studies of contour integration (Feldman, 1997; Pizlo et al., 1997) as well as visual extrapolation of contours (Singh & Fulvio, 2005; 2007).

2. Contour extrapolation and interpolation

A natural way to investigate the visual representation of contours is by examining how the visual system "fills in" the shape of contour segments that are missing in the image—for example, due to partial occlusion or camouflage (or insufficient image contrast). Shape completion is a highly under-constrained problem, a form of the problem of induction (Hume, 1748). Given any pair of inducing contour segments there are always infinitely many smooth contours that can fill-in the missing intervening portion of the shape. Because visually completed contours are, by definition, generated by the visual system (being absent in the retinal images themselves), detailed measurement of their shape provides a unique window on the shape constraints embodied in the visual processing of contours.

[3] It is important to note that, since the generative models of contours considered in this section were entirely local, these claims follow simply from *local expectations* about contour behavior.

2.1 Contour extrapolation

Perhaps the simplest context for examining visual shape completion is that of contour extrapolation: If a curved contour disappears behind an occluder, how does the visual system "expect" it will proceed behind the occluder? In other words, what shape will it take—not just in the immediate vicinity of the point of occlusion, but also further away? A precise answer to this question would serve to characterize the commonly (though often loosely) used notion of "good continuation." [4] Indeed, Wertheimer (1923) originally proposed the principle of good continuation as a way of choosing between different possible extensions of a contour segment (e.g., see his Figures 16-19). However, a mathematically precise characterization has been elusive. Some formal questions concerning the meaning of good continuation include:

i. Which geometric variables of the contour does the visual system use in extrapolating its shape, e.g., its tangent direction, curvature, rate of change of curvature, higher derivatives?
ii. How does the visual system combine the contributions of these variables to actually generate the extended shape of the extrapolated contour?

In addition, contour extrapolation is also a critical component of the general problem of shape completion—-since a visually interpolated contour must both smoothly extend each inducing contour, as well as smoothly connect the two individual extrapolants (Ullman, 1976; Fantoni & Gerbino, 2003). Therefore a full understanding of visual shape completion requires an understanding of how the visual system extrapolates each curved inducing contour.

Singh & Fulvio (2005; 2007) used an experimental method they called location-and-gradient mapping to measure the shape of visually extrapolated contours. This method obtains paired measurements of extrapolation position and orientation at multiple distances from the point of occlusion, in order to build up an extended representation of a visually-extrapolated contour. In their stimuli, a curved contour disappears behind the straight-edge of a half-disk occluder (see Figure 2a). Observers iteratively adjust the (angular) position of a short line probe on the opposite (curved) side of the occluder, and its orientation, in order to optimize the percept of smooth continuation. Measurements are taken at multiple distances from the point of occlusion, by using half-disk occluders of different sizes (see Figure 2b).

In their first study, Singh & Fulvio (2005) used arcs of circles and parabolas as inducing contours. By fitting various shape models to the extrapolation data, they found that:

i. The visual system makes systematic use of contour curvature in extrapolating contours—in other words, extrapolation curvature increases systematically with the curvature of the inducing contour. Although this result makes perfect intuitive sense, it is noteworthy that current models of shape completion (in both human and computer vision) do not use the curvature of the inducer—only its position and tangent direction at the point of occlusion. This empirical result thus underscores the need for models of shape completion to incorporate the role of inducer curvature as well.

[4] This question is of course intimately related to the generative models of contours considered in Section 1. The main difference is that previously considered models focused on where contour is likely to "next"—i.e., in the immediate vicinity of the current location—whereas the question we are now posing includes the *extended* behavior of the contour.

ii. Visually extrapolated contours are characterized by decaying curvature with increasing distance from the point of occlusion. Specifically, fits of spiral shape models (i.e., models that include both a curvature term and a rate-of-change of curvature term), to extrapolation data consistently yielded negative values for the rate-of-change of curvature.[5]

iii. The precision of subjects' visually extrapolated contours decreases systematically with the curvature of the inducing contour: the higher the inducing curvature, the less precisely the visually extrapolated contour is localized. This result is consistent with findings from contour interpolation studies using dot-sampled contours, which have also found a "cost of curvature" in human performance (Warren, Maloney, & Landy, 2002).

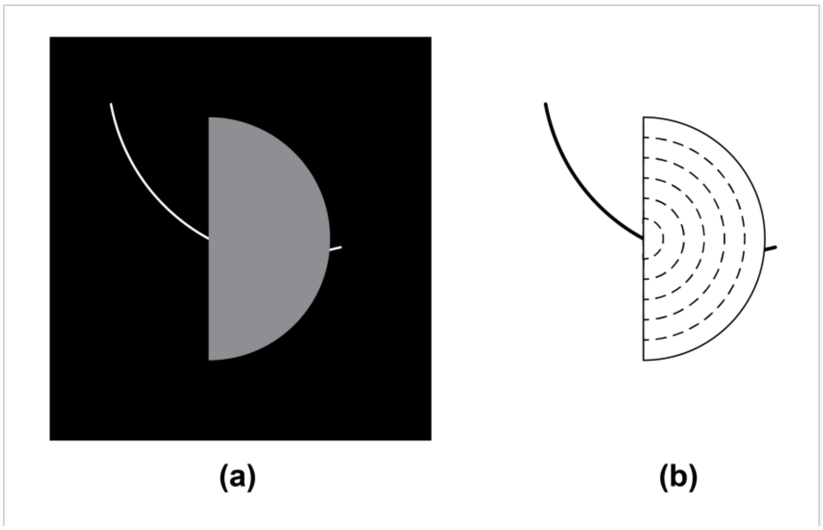

Figure 2. (a) Stimulus used by Singh & Fulvio (2005; 2007) to study the visual extrapolation of contours behind an occluder. A curved inducing contour disappears behind the straight edge of a half-disk occluder. Observers adjust the angular position as well as the orientation of a line probe around the curved edge of the occluder to optimize the percept of smooth continuation. (b) Measurements are obtained at multiple distances from the point of occlusion to build a detailed representation of an observer's visually extrapolated contour.

In a subsequent study, Singh & Fulvio (2007) tested whether observers make use of the rate-of-change of curvature of an inducing contour in visually extrapolating its shape. This study used arcs of Euler spirals as inducing contours—characterized by linearly increasing or decreasing curvature as a function of arc length (i.e., length measured along the contour)—and manipulated their rate-of-change of curvature (both in the positive and negative directions). In fitting a two-parameter Euler-spiral model to the extrapolation settings, they found no systematic relationship between the rate-of-change of curvature of the inducing contour and

[5] The decaying curvature behavior explains the (initially surprising) finding that a parabolic shape model better explained observers' extrapolation data than a circular shape model—irrespective of whether the inducing contour itself was a circular or parabolic arc (see Singh & Fulvio, 2005 for details).

the rate-of-change of curvature of the fitted Euler spiral to the extrapolation data. Thus observers appear not to take into account rate-of-change of curvature in visually extrapolating contours behind occluders. Indeed, visually extrapolated contours continued to exhibit a decaying-curvature behavior even when the inducing contour had monotonically increasing curvature as they approached the occluder. Importantly, this failure to use inducer rate-of-change of curvature was not simply due to a failure to detect it. A control experiment confirmed that observers could indeed reliably distinguish between inducing contours with monotonically increasing vs. decreasing curvature.

Taken together, these results may be viewed as providing a formal characterization of "good continuation." Specifically, they show that the visual system uses tangent direction as well as curvature—but not rate-of-change of curvature—in visually extrapolating a curved contour. Moreover, the influence of inducer curvature on visually extrapolated contours decays with distance from the point of occlusion. Singh & Fulvio (2005; 2007) modeled these characteristics using a Bayesian cue-combination model involving two probabilistically expressed constraints: a likelihood constraint to maintain the curvature of the inducing contour (i.e., a bias toward "co-circularity"; Parent & Zucker, 1989), and a prior constraint to minimize curvature (i.e., a bias toward "straightness"; e.g., Field et al., 1993; Feldman, 1997; 2001; Geisler et al., 2001; Elder & Goldberg, 2002). Both constraints were expressed as probability distributions on curvature. The prior was expressed as a Gaussian distribution centered on 0 curvature with fixed variance; whereas the likelihood was centered on the estimated inducer curvature at the point of occlusion, with a (Weber-like) linearly increasing standard deviation with distance from the point of occlusion. Near the point of occlusion, the likelihood is very precise (low variance) and thus tends to dominate the prior.[6] With increasing distance from the point of occlusion, however, the likelihood becomes less reliable (larger variance), and so gradually the prior comes to dominate the likelihood. This shift in relative reliabilities leads to the decaying curvature behavior (see Singh & Fulvio, 2007, for details).

2.2 Contour interpolation

Fulvio, Singh & Maloney (2008) extended the location-and-gradient mapping method to study contour interpolation. Their stimulus displays contained a contour whose middle portion was occluded by a rectangular surface. On each trial, a vertical interpolation window was opened at one of six possible locations, through which a short linear probe was visible (see Figure 3a). Observers iteratively adjusted the location (height) and orientation of the line probe in order to optimize the percept of smooth continuation of a single contour behind the occluder. The perceived interpolated contours were thus mapped out by taking measurements at six evenly spaced locations along the width of the occlusion region. The experiments manipulated the geometry of the two inducing segments—specifically, the turning angle between them (Figure 3b) and their relative vertical offset (Figure 3c).

A basic question was: for a given pair of inducing contours, are observers' settings of position and orientation through the six interpolation windows globally consistent—i.e., consistent with a single, stable, smooth interpolating contour. Using two measures of global consistency—a

[6] Under the assumption of Gaussian distributions for the prior and likelihood, the Bayesian posterior is also a Gaussian distribution whose mean is a weighted average of the prior mean and likelihood mean, with the relative weights inversely proportional to their respective variances (see, e.g., Box & Tiao, 1992).

parametric one and a non-parametric one—Fulvio et al. (2008) found that although increasing the turning angle between inducers adversely affected the precision of interpolation settings, it did not adversely affect their internal consistency. By contrast, increasing the relative offset between the two inducing contours did disrupt the internal consistency of observers' interpolation settings. In other words, observers made their settings using simple heuristics (they were largely influenced by the closest inducing contour), and their local settings of height and orientation at various locations no longer "hung together" into any actual extended contour. A natural way to understand this difference is that increasing the relative offset between inducer pairs leads eventually to a geometric context where the interpolating contour must be inflected—i.e., contain a point of inflection (or change in the sign of curvature) somewhere along its path—which is a factor that is known to disrupt visual completion (Takeichi et al., 1995; Singh & Hoffman, 1999). On the other hand, simply increasing the turning angle between the two inducers does not necessitate inflected interpolating contours—it only requires interpolating contours with greater curvature in a single direction.

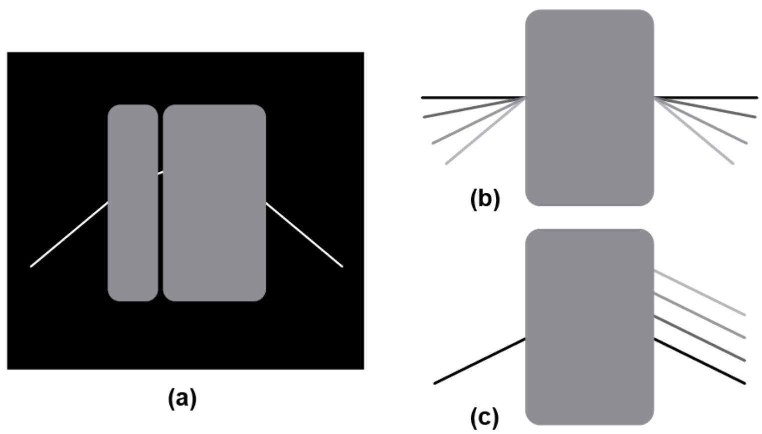

Figure 3. (a) Stimulus used by Fulvio, Singh & Maloney (2008, 2009) to study contour interpolation. For a given pair of inducing edges, an interpolation window is opened at one of six possible locations along the width of the occluder. Observers adjust the height as well as the orientation of a line probe visible through the interpolation window in order to optimize the percept of smooth interpolation. The inducer geometry was manipulated by varying the turning angle (show in (b)) and the relative offset (shown in (c)) between the two inducers.

These two factors—turning angle and relative offset between inducers—are often combined conjunctively to define the strength of grouping between pairs of inducing edges. For example, Kellman & Shipley's (1991) definition of edge relatability requires that both the relative offset between inducers, as well as the turning angle between them, be within specific ranges in order for them to be considered "relatable." This conjunctive combination, however, ignores the qualitatively different effects that these two factors have on contour interpolation. Specifically, although both factors lead to an increase in imprecision, only relative offset leads to a failure of internal consistency. In a subsequent study, Fulvio et al. (2009) developed a purely experimental criterion to test for internal consistency of interpolation measurements—one that relied solely on observers' own interpolation performance rather than on any experimenter-defined measures. The results independently verified and extended their earlier findings.

3. Part-based representations of shape

A great deal of evidence—both psychophysical (see below) and physiological (e.g., Pasupathy & Connor, 2002)—indicates that the human visual system represents contours and shapes in a piecewise manner. In other words, it segments contours and shapes into simpler "parts" and organizes shape representation using these parts and their spatial relationships. Far from being arbitrary subsets, these perceptual parts are highly systematic, and segmented using predictable geometric "rules." Moreover, these segmented parts tend to correspond, in high-level vision, to psychologically meaningful sub-units of objects (such as head, leg, branch, etc.) that are highly relevant to a number of cognitive processes, including categorization, naming, and object recognition.

Although in Attneave's (1954) usage, the phrase "maxima of curvature" along a contour does not distinguish between positive (convex) and negative (concave) curvature, the sign of curvature actually plays a fundamental role in modern theories of shape representation—and especially in theories of part segmentation. Once one treats curvature as a signed quantity (which can be done whenever the distinction between convex and concave is well-defined), one can differentiate between positive maxima of curvature (marked by **M+** in Figure 4a) and negative minima of curvature (marked by **m-** in Figure 4a). Both of these extrema types have locally maximal magnitude of curvature, and are hence "maxima of curvature" by Attneave's nomenclature. However, by definition, positive maxima lie in convex segments of a shape's bounding contour, whereas negative minima lie in concave segments. Apart from these two extrema types, another important class of points is defined by inflections, which are zero crossings of curvature—i.e., points where curvature crosses from positive (convex) to negative (concave), or vice versa (marked by **o** in Figure 4a).

The distinction between positive maxima and negative minima of curvature is critical for part segmentation—where negative minima of curvature play a special role. According to Hoffman & Richards' (1984) "minima rule," the visual system uses negative minima of curvature to segment shapes into parts. This rule is motivated by the principle of transversality, according to which when two smooth objects are joined to form a composite object, their intersection generically produces a concave crease (i.e., a discontinuity in the tangent plane of the composite surface; see Figure 4b). And a concave crease is simply an extreme—i.e., "sharp"—form of a negative minimum of curvature. (More precisely, a generic application of smoothing to a concave crease yields a smooth negative minimum.) Similarly, when a new branch grows out of a trunk (or a limb out of an embryo), negative minima of curvature are created between the sprouting branch and the trunk (see Figure 4c; Leyton 1989). Hence, when faced with a complex object with unknown part structure, it is a reasonable strategy for the visual system to use the presence of negative minima of curvature as a cue to identifying separate parts.

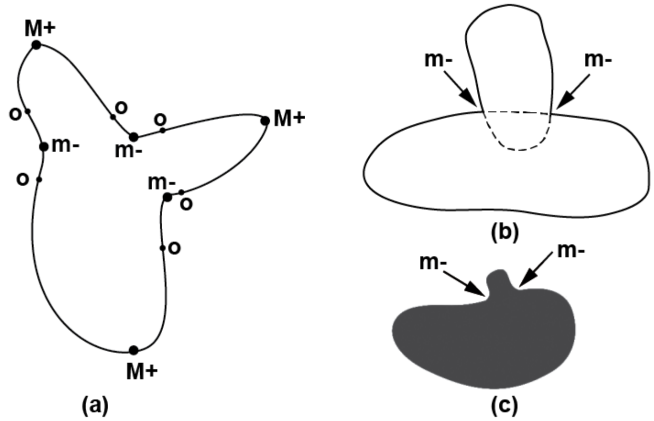

Figure 4. (a) Illustrating different types of curvature-based features along the outline of a shape: Positive maxima of curvature (marked by **M+**), negative minima of curvature (marked by **m-**), and inflection points (marked by **o**). (b) Motivation behind the minima rule: Joining two smooth objects generically produces negative minima of curvature on the composite object. (c) Similarly, when a branch grows out of a trunk (or a limb out of an embryo), negative minima are created at the loci of protrusion.

A great deal of psychophysical evidence indicates that negative minima of curvature do indeed play an important role in visually segmenting shapes into parts. For example, when subjects are asked to draw cuts on line drawings of various objects to demarcate their natural parts, a large proportion of their cuts pass through or near negative minima of curvature (Siddiqi et al., 1996; De Winter & Wagemans, 2006). Similar results have also been obtained with 3D models of objects (Chen, Golivinskiy & Funkhouser, 2009). Furthermore, even when unfamiliar, randomly-generated, shapes are used (hence lacking any high-level cues from recognition or category knowledge), and subjects are simply asked to indicate whether or not a given contour segment belongs to a particular shape (i.e., in a performance-based task where the instructions to participants involve no mention of "parts"), their identification performance is substantially better for segments delineated by negative minima of curvature than for those delineated by other extrema types (Cohen & Singh, 2007). This result indicates that part segmentation is a relatively low-level geometry-driven process that operates automatically without relying on familiarity with the shape, or any task requirement involving naming or recognition.[7]

Part segmentation using negative minima of curvature has been shown to explain a number of visual phenomena, including the perception of figure and ground (Baylis & Driver, 1994; 1995; Hoffman & Singh, 1997), the perception of shape similarity (Hoffman & Richards, 1984; Bertamini & Farrant, 2005; Vandekerckhove, Panis & Wagemans, 2008), object recognition in contour-deleted images (Biederman, 1987; Biederman & Cooper, 1991); perception of

[7] This does not mean, of course, that high-level cognitive factors do not also exert an influence when present; they clearly do (see, e.g., De Winter & Wagemans, 2006). The point is simply that cognitive factors are not *necessary* for part segmentation; low-level geometry-driven mechanisms of part segmentation can and do operate in their absence.

transparency (Singh & Hoffman, 1998), visual search for shapes (Wolfe & Bennett, 1997; Hulleman, te Winkel & Boselie, 2000; Xu & Singh, 2001), the visual estimation of the "center" of a two-part shape (Denisova, Singh & Kowler, 2006); the visual estimation of the orientation of a two-part shape (Cohen & Singh, 2006), and the allocation of visual attention to multi-part objects (Vecera, Behrmann & Filapek, 2001; Barenholtz & Feldman, 2003).

Although the minima rule provides an important cue for part segmentation, it is not sufficient to divide a shape into parts—which of course requires segmenting the interior region of a shape, not simply its bounding contour. Specifically, although the minima rule provides a number of candidate part boundaries (namely, the negative minima of curvature), it does not indicate how these boundaries should be paired to form part cuts that segment the shape. Furthermore, even in shapes containing exactly two negative minima, simply connecting these two minima does not necessarily yield intuitive part segmentations (see e.g., Singh, Seyranian & Hoffman, 1999; Singh & Hoffman, 2001 for examples). The basic limitation of the minima rule stems from the fact that localizing negative minima of curvature involves only the local geometry of the bounding contour of the shape, but not the nonlocal geometry its interior region (see Section 4 for more on this important distinction). Because of the contributions of such nonlocal region-based factors, it is possible to have negative minima on a shape that do not correspond to perceptual part boundaries (Figure 5a) and, conversely, to have perceptual part boundaries that do not correspond to negative minima (Figure 5b).

In order to address such limitations, researchers have proposed a number of additional geometric factors for segmenting objects into parts: limbs and necks (Siddiqi, Tresness & Kimia 1996), convexity (Latecki & Lakamper, 1999; Rosin, 2000), a preference for shorter cuts (Singh et al., 1999), local symmetry, good continuation (Singh & Hoffman, 2001), as well as cognitive factors based on object knowledge (De Winter & Wagemans, 2006). And each of these factors has indeed been shown to play a role in part segmentation. However, with a large number of such factors (in addition to the minima rule), it becomes increasingly difficult to model the various complex interactions between them—the way in which they co-operate and compete with each other in various geometric contexts—and therefore to have a unifying theory of part segmentation.

A different approach to part segmentation is to use an axial, or skeleton-based, representation of the interior region of a shape in order to segment it into parts. Specifically, each axial branch of the shape skeleton can be used to identify a natural part of the shape (see Figure 5c)—assuming, of course, that the skeleton-computation procedure can yield a one-to-one correspondence between parts and axial branches. The desirability of such a correspondence was in fact articulated in Blum's original papers that introduced his Medial-Axis Transform (MAT) as a representation of animal and plant morphology (e.g., Blum, 1973).[8] However, as recognized subsequently by Blum & Nagel (1979; see their Figure 2), the MAT does not achieve this one-to-one correspondence. Although modern techniques for computing the medial axis and related transforms have become increasingly sophisticated, they nevertheless largely inherit the intrinsic limitations of the MAT—which follow from the basic conception of skeleton computation as a deterministic process involving the application of a fixed geometric "transform" to any given shape. Specifically, a geometric-transform approach does not attempt to separate

[8] In the MAT conception, a shape is viewed as the union of maximally inscribed circles, and its skeleton—the MAT—is taken to be the locus of the centers of these circles.

the shape "signal" from any contributions of noise. Every feature along the contour is effectively treated as being "intrinsic" to the shape. One consequence of this is a high degree of sensitivity of the skeleton to noise, such that the smallest perturbation to the contour can dramatically alter the branching topology of the shape skeleton.

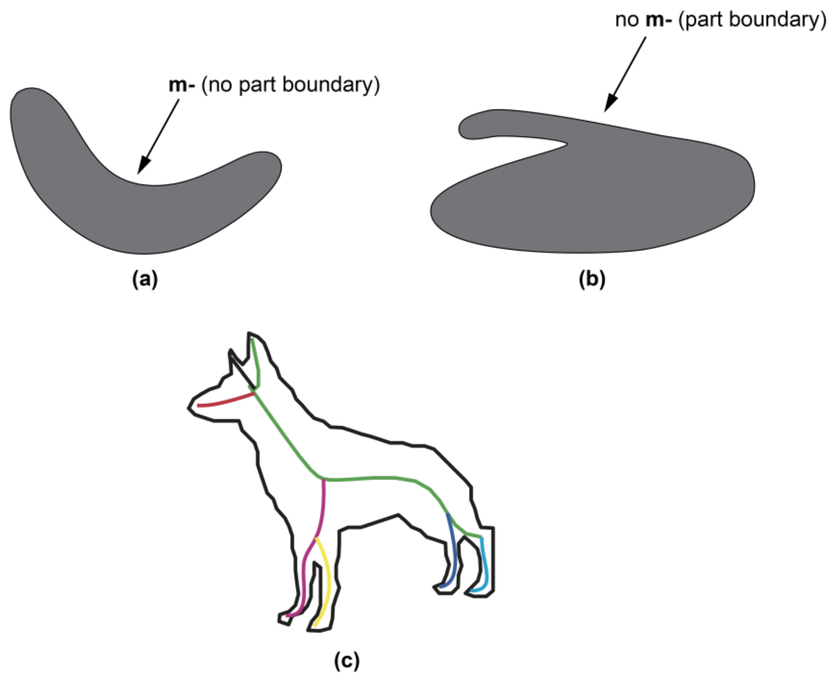

Figure 5. Two examples of failure of the minima rule: (a) A negative minimum that does not correspond to a part boundary; and (b) a part boundary that does not correspond to a negative minimum. These failures arise because the minima rule uses only local contour geometry, not region-based geometry. (c) A different approach to part segmentation involves establishing a one-to-one correspondence between axial branches are parts. Such a correspondence is achieved by a Bayesian approach to skeleton computation (Feldman & Singh, 2006).

In order to address these concerns, Feldman & Singh (2006) used an inverse-probability approach to estimate the skeleton that "best explains" a given shape. The key idea in this approach is to treat object shapes as resulting from a combination of generative factors and noise. The skeletal shape representation must then model the generative (or "intrinsic") factors, while factoring out the noise. Specifically, shapes are assumed to "grow" from a skeleton via a stochastic generative process. The estimated skeleton of a given shape is then one's best inference of the skeleton that generated it. Skeletons with more branches, and more highly curved branches, can of course provide a better fit to the shape (i.e., lead to a higher likelihood), but they are also penalized for their added complexity (i.e., they have a lower prior). Thus one's "best" estimate of the skeleton involves a Bayesian tradeoff between fit to the shape, and the complexity of the skeleton.

This tradeoff leads to a pruning criterion for "spurious" branches of the shape skeleton: a candidate axial branch is included in the final shape skeleton only if it improves the fit to the shape sufficiently to warrant the increase in skeletal complexity that it entails. More precisely, the posterior of the skeleton that includes the test branch must be larger than the posterior of the skeleton that excludes it (recall that the posterior includes both the contribution of the fit to the shape, via the likelihood term, as well as of skeleton complexity, via the prior.) Axial branches that do not meet this criterion are effectively treated as "noise" and pruned. As a result, this probabilistic computation is able to establish a one-to-one correspondence between axial branches and perceptual parts (see Figure 5c for an example). Importantly, it can predict both the successes of the minima rule (cases where negative minima are perceived as part boundaries) and its failures (cases where negative minima are not perceived as part boundaries, or where part boundaries do not correspond to negative minima; recall Figures 5a and 5b)—despite the fact that in this approach contour curvature is never explicitly computed. Thus, it would yield a single axial branch for the curved shape in Figure 5a; but a skeleton with two axial branches for the shape in Figure 5b. Indeed, the contributions of other known factors influencing part segmentation can all be understood in terms of this more fundamental process of probabilistic estimation of the shape skeleton, indicating that this may provide a unifying theory of part segmentation. See Singh, Feldman & Froyen (in preparation) and Feldman, Singh, Briscoe, Froyen, Kim & Wilder (2013) for more on this probabilistic approach to skeletons and parts, and its application to various visual problems.

4. Interactions between contour and region geometry

The Gestaltists noted early on that a closed contour is perceptually much more than an open one (Koffka, 1935). And this claim has been corroborated in a number of experimental contexts (e.g., Elder & Zucker, 1993; Kovacs & Julesz, 1993; Garrigan, 2012). However, because closed contours automatically define an enclosed region, it is less clear whether this advantage of closure obtains at the level of contour geometry (see Tversky, Geisler & Perry, 2004), or at the level of region-based geometry, i.e., the geometry of the region enclosed by the contour.

We have seen in the context of part segmentation that there is more to the representation of a shape than simply the geometry of its bounding contour. To motivate the distinction between contour geometry and region (or surface) geometry further, consider the simple shape shown in Figure 6a. This shape may be conceptualized in two different ways:

i. It could be viewed as a rubber band lying on a table (the "rubber-band representation"). Mathematically, we would define it as a closed one-dimensional contour embedded in two-dimensional space. In this case, a natural way to represent its geometry would be in terms of some contour property—say, curvature—expressed as a function of arc length (resulting in a curvature plot such as in Figure 6b). The relevant notions of distance and neighborhood relations would then also be defined along the contour. As a result, although points **A** and **B** on the shape are close to each other in the Euclidean plane, they would not be considered "neighboring" points because they are quite far from each other when distances are measured along the contour.

ii. Alternatively, it could be viewed as a piece of cardboard cut out into a particular shape (the "cardboard-cutout representation"). Mathematically, we may define it as a connected and compact two-dimensional subset of the Euclidean plane (namely, the region enclosed by the contour). Under this conceptualization, points **A** and **B** on the shape would indeed be considered quite close to each other (because the intervening region is now also part of the shape).

The distinction between region-based and contour-based notions of shape has a number of other implications as well. In Figure 6c, for example, the two highlighted sections of the contour belong to the same "bend" in the shape. A purely contour-based representation, however, would have difficulty in explicitly representing this fact. In the curvature plot in Figure 6d, for instance, the two contour sections do not appear to be related in any obvious way. What a contour-based representation misses here is the locally parallel structure of the two highlighted contour segments. It is clear that such structure can be extracted only by examining relationships across (i.e., on "opposite" sides of) the shape—not just along the contour. For the same reason, bilateral symmetry or local symmetry in shapes is relatively easy to capture using region-based representations, but difficult using purely contour-based representations. As an example, even though the two shapes shown in Figure 7 have very similar curvature profiles, their global region-based geometries are entirely different (Sebastian & Kimia, 2005).

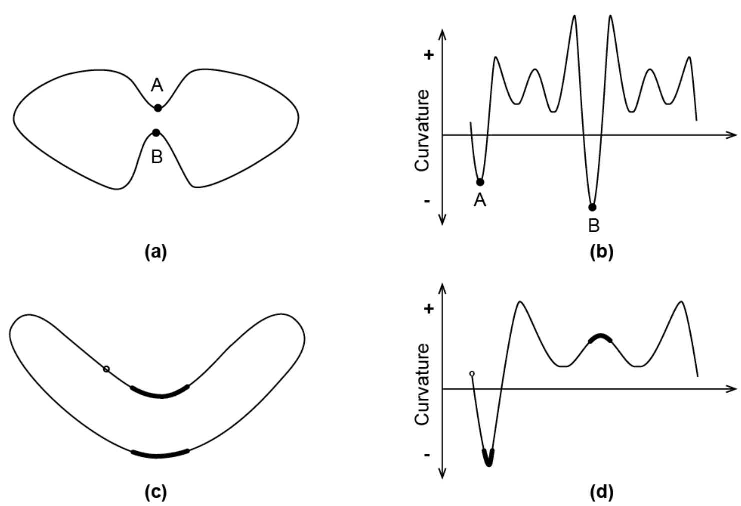

Figure 6. Illustrating the limitations of a contour-based representation of shape. (a) Although the two points **A** and **B** are very close to each other on the shape, they are very distant on the curvature plot of its bounding contour, as shown in (b). (c) Similarly, although the two highlighted sections of the contour belong to the same "bend" in the shape, this fact is not reflected in any obvious way from the curvature plot in (d).

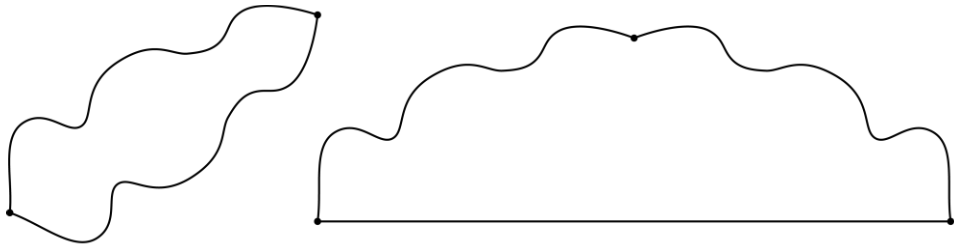

Figure 7. Although the two shapes have similar curvature profiles—differing only in the presence of a zero-curvature segment in the shape on the right—their region-based geometries are entirely different. Example based on Sebastian & Kimia (2005).

We should note that, in the examples above, we assumed that "material" surface was on the inside of the closed contour—not an unreasonable assumption for closed contours if we know we are viewing solid, bounded, objects (the alternative would be an extended surface containing a shaped hole). In the general case, however, the visual system faces the problem of border-ownership or figure-ground assignment—determining whether the material object or surface lies on one side of the contour or the other—a problem that is particularly acute when only a small portion of an object's outline is visible. An interesting interaction occurs between contour geometry and region-based geometry in solving this problem, such that the side with the "simpler" region-based description tends to be assigned figural status. In more formal terms, the relevant geometric factors have been characterized in terms of part salience (Hoffman & Singh, 1997) and stronger axiality (Froyen, Feldman & Singh, 2011).

A natural way to capture region-based geometry is in terms of skeletal, or axial, representations (introduced briefly in Section 3)—compact "stick-figure" representations that capture essential aspects of its morphology (see, e.g., Kimia, 2003). A well-known figure by Marr & Nishihara (1978) shows 3D models of various animals made out pipe cleaners. A striking aspect of these models is how easily they are recognized as specific animals, despite the absence of surface geometry—or indeed any surface characteristics. The demonstration suggests that the axial information preserved in these pipe-cleaner models is an important component of human shape representation. It should be borne in mind, however, that a skeletal representation actually includes not just an estimate of the shape's axes (which are shown in Marr & Nishihara's pipe-cleaner models), but also an estimate of the shape's "width" at each point on each axis (which is not). In Blum's MAT, for instance, this local "width" is captured by the size of the maximally inscribed circle at any given point. In Feldman & Singh's (2006) Bayesian skeleton model, it is approximately twice the length of the "ribs" along which the shape is assumed to have "grown" from the axis. Each such measure of local width of the shape implicitly defines a point-to-point correspondence across the shape. In other words, it specifies, for any given point on the shape's bounding contour, which point on the "opposite" side of the shape is locally symmetric to it.[9]

[9] One way to think about local symmetry is as follows: Imagine placing a mirror at a point along the shape's axis, with its orientation matching the local orientation of the axis. If the axis is defined appropriately, the mirror will reflect the tangent of the contour on one side of the shape to the tangent of the contour on the opposite side of the shape (Leyton, 1989).

What are the perceptual implications of the difference between contour-based geometry and region-based geometry? Consider the local contour segment in Figure 8a, shown through an aperture. The same contour segment could belong to shapes with very different region-based geometries. First, the contour segment could correspond either to a convex protuberance on the shape, or to a concave indentation (Figures 8b vs. 8c). This distinction is based simply on a figure-ground reversal (or change in border ownership)—whether the shape lies either on one, or the other, side of the contour. And this has been shown to be an important factor in predicting perceptual grouping in the context of both amodal (Liu et al., 1999) and modal (Kogo et al., 2010) completion.

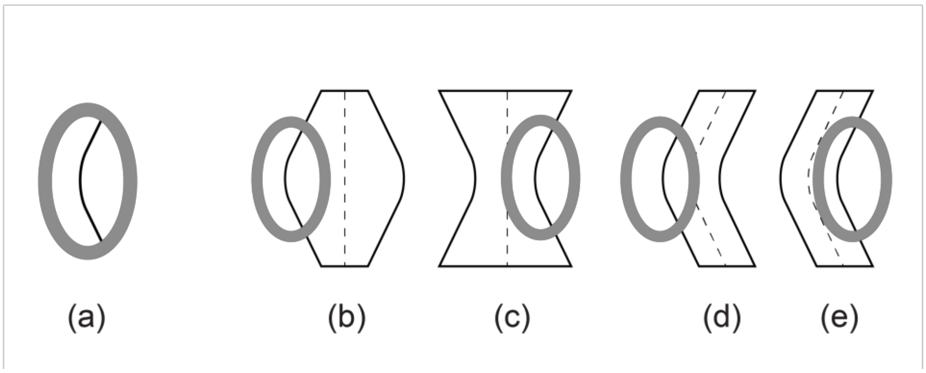

Figure 8. Illustrating the distinction between contour and region (or surface) geometry. The same contour segment, visible through an aperture in (a), could belong to surfaces with very different geometries. First, the contour segment could correspond to a protuberance on the shape, as in (b), or to an indentation, as in (c). Second, the curvature of the contour could arise due to variation in the width of the shape about a straight axis (as in (b) and (c)), or due to curvature of the axis itself, with the local width function being constant (as in (d) and (e)).

The second distinction we consider, however, does not depend on a figure-ground reversal: assuming a locally convex region (say), the curvature on the contour could arise either from variation in the width of the shape about a straight axis (as in Figures 8b and 8c), or from curvature of the axis itself, with the local width of the shape being constant (Figures 8d and 8e). It is clear that these two cases actually represent two extremes of a continuum—where all of the contour curvature can be attributed entirely to either the width function alone, or to axis curvature alone. A continuous family of intermediate cases is of course possible—where the contour's curvature arises partly due to the curvature of the shape's axis, and partly due to variations in the shape's width (Siddiqi et al., 2001; Fulvio & Singh, 2006).

In order to examine the perceptual consequences of such region-based differences in shape, Fulvio & Singh (2006) examined visual shape interpolation in stereoscopic illusory-contour displays. Their displays varied systematically in their region-based geometry, while preserving the contour-based geometry of the inducing edges (see Figure 9). Using two different experimental methods, they probed the perceived shape of the illusory contours in the "missing" region. The results exhibited large influences of region-based geometry on perceived illusory-contour shape. First, illusory contours enclosing locally concave shapes were found to be systematically more angular (closer to the intersection point of the linear extrapolations of the two inducers) than those enclosing locally convex shapes. This influence of local convexity is consistent with results obtained with partly occluded shapes (Fantoni et al., 2005). Beyond the influence of local sign of curvature, however, this influence of local convexity also exhibited an interaction with two skeleton-based variables: shape width and axis curvature. Specifically, the influence of local convexity on illusory-contour shape was found to be: (i) greater for narrower shapes than for wider ones; and (ii) greater for shapes with a straight axis and symmetric contours ("diamonds" and "bowties"; Figure 8b and 8c) than for shapes with a curved axis and locally parallel contours ("bending tubes"; see Figures 8d and 8e). These results indicate that, even at the level of illusory "contours," an important role is played by nonlocal region-based geometry involving skeleton-based parameters.

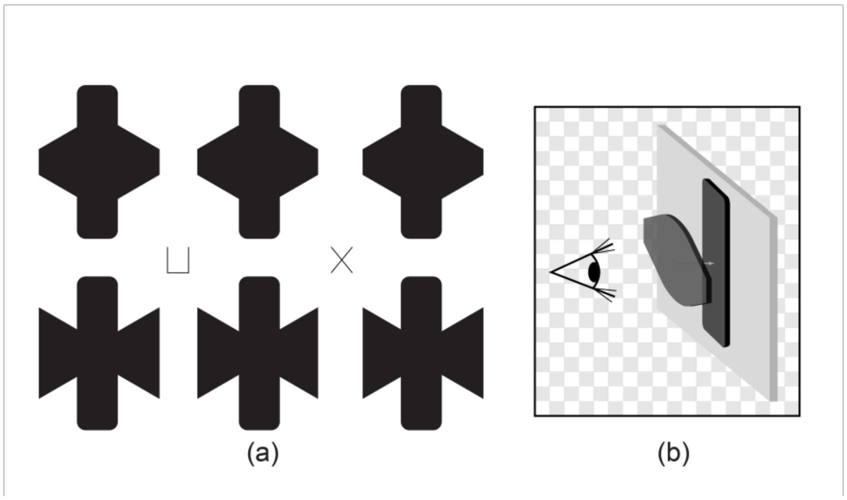

Figure 9. (a) Stereoscopic stimuli used by Fulvio & Singh (2006) to study the influence of region-based geometry on illusory-contour shape. In these stimuli, region-based geometry was manipulated while keeping local contour geometry fixed (as in Figure 8). A schematic of the binocular percept is shown in (b). The results showed significant differences in perceived illusory-contour shape as a function of region-based geometry.

The influence of region-based geometry manifests itself in object recognition and classification as well. In comparing the recognition performance of contour and region-based models, Sebastian & Kimia (2005) compared the shape matching performance of two algorithms—one based on matching their bounding contours, the other based on matching axis-based graphs derived from them. They found that when small variations were introduced on the shapes (e.g., involving partial occlusion, rearrangement of parts, or addition or deletion of a part), the contour-based matching scheme produced many spurious matches, leading to a substantial deterioration in performance. By contrast, the axis-based matching scheme was highly robust to such variations. They concluded that, even though axis-based representations are more complex and take more time to compute, the additional time and effort required to compute them are well worth it.

Do human observers make use of parameters of the shape skeleton in classifying shapes? Different classes of shapes—e.g., animals and leaves—differ not only in their means along various skeleton-based parameters (e.g., number of branches, axis curvature, etc.), but also in their distributional forms. For example, the distribution of number of branches tends to be Gaussian for animals with a mean of around 5 (reflecting the typical number of body parts in an animal body plan), whereas the distribution tends to be exponential for leaves (consistent with a recursively branching process); see Figure 10a. Do human subjects rely on such statistical differences in skeletal parameters when performing shape classification? Wilder, Feldman & Singh (2011) used morphed shapes created by combining animal and leaf shapes in different proportions (e.g., 60% animal and 40% leaf; see Figure 10b). Subjects indicated whether each shape looked more like an animal, or more like a leaf. (The morphing proportions ranged between 30% and 70% so the shapes were typically not recognizable as any particular animal or leaf). They then compared subjects' performance with that of a naive Bayesian classifier based on a small number of skeletal parameters, and found a close match between the two. By contrast, classification based only on contour-based variables (such as contour curvature) and other traditional shape measures (such as compactness and aspect ratio) did not provide good predictors of human classification performance. These comparisons provide strong evidence for the use of a skeleton-based representation of shape by the human visual system. More recent work also provides evidence for the role of region-based representation of shape in contour detection tasks, i.e., detecting a closed contour in background noise (Wilder, Singh & Feldman, 2013).

5. Conclusions

Contours constitute an essential source of information about shape and, along contours, points with the greatest magnitude of curvature tend to be most informative. This concentration of information is closely tied to generative models of contours assumed by the visual system—i.e., its internal models about how contours tend to be generated (and hence its expectations about how contours tend to behave locally). Therefore visual expectations about contour continuity (“good continuation”) and the information content of contours are naturally viewed as two sides of the same coin. In going from open to closed contours—such as the outlines of objects—the influence of sign of curvature (convex vs. concave) becomes critical, with concave sections of a contour carrying more information, and playing a special role in part segmentation. The visual system represents complex shapes by automatically segmenting them into simpler parts—“simpler” because these parts are closer to being convex (they contain less negative curvature). One type of curvature extrema—negative minima of curvature—provides a particularly important cue for part segmentation. However, sign of curvature (local convexity) and curvature extrema are entirely contour-based notions, and this fact likely explains why the minima rule cannot fully predict part segmentation. The visual system employs not only a contour-based representation of shape, but also a region-based one—namely, a representation of the interior region enclosed by the contour—making explicit properties such as the local width of the shape, the curvature of its axis, and more generally, locally parallel and symmetric structure. Psychophysical results from a variety of domains—shape classification, amodal and modal grouping, visual shape completion—provide clear evidence for the representation of region geometry based on skeleton or axis models. Even at the level of so-called “illusory contours,” nonlocal region-based geometry exerts a strong influence.

We conclude that, as far as the human visual representation of shape is concerned, contour geometry cannot ultimately be viewed in isolation, but must be considered in tandem with region-based geometry.

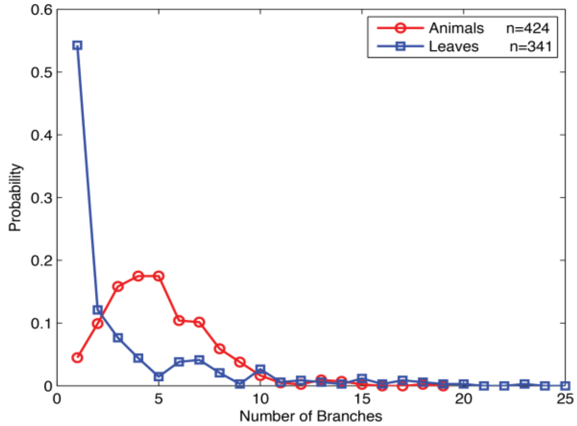

(a)

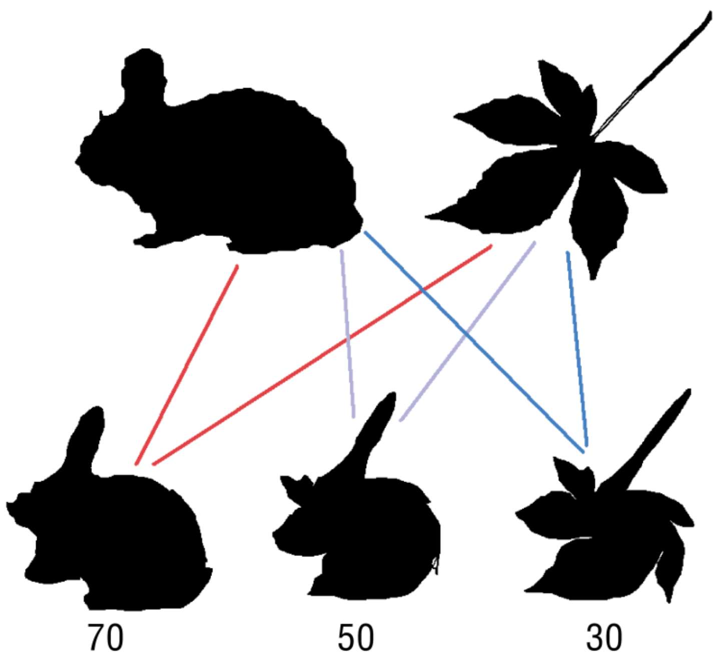

(b)

Figure 10. Different categories of shape, such as animals and leaves, differ in the statistics of various skeleton-based parameters. (a) shows the distribution of number of axial branches computed from databases of animal and leaf shapes. Note that the two categories differ both in the mean, as well as the distributional form, of this variable. (b) To address the question of whether human observers rely on skeleton-based statistics to classify shapes, Wilder, Feldman & Singh (2011) created morphed shapes by mixing animals and leaves in different proportions. Subjects were asked whether each morphed shape looked “more like” an animal or leaf. The results showed that a naive Bayesian classifier based on the distribution of a small number of axis-based parameters provided an excellent predictor of human shape classification.

6. 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(15), 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. & 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(6):1323–1342.

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.

Biederman, I., & Ju, G. (1988). Surface vs. Edge-Based Determinants of Visual Recognition. Cognitive Psychology, 20, 38–64

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.

Box, G. E. P., & Tiao, G. C. (1992). Bayesian inference in statistical analysis. New York: Wiley.

Chen, X., Golovinskiy, A. & Funkhouser, T. A. (2009). A benchmark for 3D mesh segmentation. ACM Transactions on Graphics, 28(3).

Clottes, J. (2003). Chauvet Cave: The Art of Earliest Times. (Translated by Paul G. Bahn.) University of Utah Press.

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):4, 805–821.

Cohen, E. H., & Singh, M. (2007). Geometric determinants of shape segmentation: Tests using segment identification. Vision Research, 47, 2825–2840.

Cole, F., Sanik, K., DeCarlo, D., Finkelstein, A., Funkhouser, T., Rusinkiewicz, S. & Singh, M. (2009). How well do line drawings depict shape? In ACM Transactions on Graphics (Proc. SIGGRAPH), volume 28, 2009.

De Winter, J., & Wagemans, J. (2006). Segmentation of object outlines into parts: A large-scale integrative study. Cognition, 25, 275–325.

De Winter, J., & Wagemans, J. (2008a). The awakening of Attneave's sleeping cat: Identification of everyday objects on the basis of straight-line versions of outlines. Perception, 37, 245–270.

De Winter, J., & Wagemans, J. (2008b). Perceptual saliency of points along the contour of everyday objects: A large-scale study. Perception & Psychophysics, 70(1), 50–64.

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

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.

Elder, J. H. & Zucker, S. W. (1993). Contour closure and the perception of shape. Vision Research, 33(7), 981–991.

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. (1997). Curvilinearity, covariance, and regularity in perceptual groups. Vision Research, 37(20), 2835–2848.

Feldman, J. (2001). Bayesian contour integration. Perception & Psychophysics, 63(7), 1171–1182.

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, 103(47):18014–18019.

Feldman, J., Singh, M., Briscoe, E., Froyen, V., Kim, S. & Wilder, J. (2013). An integrated Bayesian approach to shape representation and perceptual organization. In: S. Dickinson & Z. Pizlo (Eds.), Shape Perception in Human and Computer Vision: An Interdisciplinary Perspective (pp. 55-70). Springer Verlag.

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

Froyen, V., Feldman, J., & Singh, M. (2010). A Bayesian framework for figure-ground interpretation. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. Zemel, & A. Culotta (Eds.), Advances in neural information processing systems, 23 (pp. 631–639).

Fulvio, J.M. & Singh, M. (2006). Surface geometry influences the shape of illusory contours. Acta Psychologica, 123, 20-40.

Fulvio, J.M., Singh, M., & Maloney, L.T. (2008). Precision and consistency of contour interpolation. Vision Research, 48, 831–849.

Fulvio, J.M., Singh, M., & Maloney, L.T. (2009). An experimental criterion for consistency in interpolation of partially-occluded contours. Journal of Vision, 9(4):5, 1-19.

Garrigan, P. (2012). The effect of contour closure on shape recognition. Perception, 41(2), 221-235.

Geisler, W. S., & Perry, J. S. (2009). Contour statistics in natural images: Grouping across occlusions. Visual Neuroscience, 26, 109–121.

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

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

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

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

Hume, D. (1748/1993). An enquiry concerning human understanding. Indianapolis, IN: Hackett Publishing.

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

Kennedy, J. M., Domander, R. (1985). Shape and contour: The points of maximum change are least useful for recognition. Perception, 14, 367–370

Kimia, B. (2003). On the role of medial geometry in human vision. Journal of Physiology Paris, 97:155–190.

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

Koenderink, J. J. (1984). What does the occluding contour tell us about solid shape? Perception, 13, 321–330.

Koffka, K. (1935). Principles of Gestalt Psychology. New York: Harcourt, Brace and World.

Kogo, N., Strecha, C., Van Gool, L., Wagemans, J. (2010). Surface construction by a 2-D differentiation-integration process: A neurocomputational model for perceived border ownership, depth, and lightness in Kanizsa figures. Psychological Review, 117(2), 406-439.

Kovacs, I. & Julesz, B. (1993). A closed curve is much more than an incomplete one: Effect of closure in figure-ground segmentation. Proceedings of the National Academy of Sciences, 90, 7495–7497.

Latecki, L., & Lakamper, R. (1999). Convexity rule for shape decomposition based on discrete contour evolution. Computer Vision and Image Understanding, 73, 441–454.

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

Panis, S., de Winter, J., Vandekerckhove, J., & Wagemans, J. (2008). Identification of everyday objects on the basis of fragmented versions of outlines. Perception, 37, 271–289.

Parent, P. & Zucker, S.W. (1989). Trace inference, curvature consistency, and curve detection. IEEE Transactions on pattern analysis and machine intelligence, 2(8), 823-839.

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

Pizlo, Z., Salach-Goyska, M., & Rosenfeld, A. (1997). Curve detection in a noisy image. Vision Research, 37(9), 1217–1241.

Richards, W., Dawson, B., & Whittington, D. (1986). Encoding contour shape by curvature extrema. Journal of the Optical Society of America A, 3, 1483–1491.

Rosin, P. L. (2000). Shape partitioning by convexity. IEEE Trans. Systems, Man, and Cybernetics, Part A, 30: 202–210.

Sebastian, T. and Kimia, B. (2005). Curves vs. skeletons in object recognition. Signal Processing, 85 (2), 247–263.

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

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

Singh, M., & Feldman, J. (2012). Principles of contour information: a response to Lim and Leek (2012). Psychological Review, 119(3), 678–683.

Singh, M., Feldman, J. & Froyen, V. (in preparation). Unifying parts and skeletons: a Bayesian approach to part segmentation. To appear in the Handbook of Computational Perceptual Organization, Oxford University Press.

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., & Fulvio, J. M. (2007). Bayesian contour extrapolation: Geometric determinants of good continuation. Vision Research, 47, 783–798.

Singh, M. & Hoffman, D. D. (1998). Part boundaries alter the perception of transparency. Psychological Science, 9, 370-378.

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). New York: Elsevier.

Singh, M., Seyranian, G. D., & Hoffman, D. D. (1999). Parsing silhouettes: the short-cut rule. Perception & Psychophysics, 61(4), 636–660.

Takeichi, H., Nakazawa, H., Murakami, I., & Shimojo, S. (1995). The theory of the curvature-constraint line for amodal completion. Perception, 24, 373–389.

Tversky, T., Geisler, W., & Perry, J. (2004). Contour grouping: closure effects are explained by good continuation and proximity. Vision Research, 44(24):2769--2777.

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

Vandekerckhove, J., Panis, S., & Wagemans, J. (2008). The concavity effect is a compound of local and global effects. Perception & Psychophysics, 69, 1253–1260.

Vecera, S.P, Behrmann, M., & Filapek, J.C. (2001). Attending to the parts of a single object: Part-based selection limitations. Perception & Psychophysics, 63, 308–321.

Walther, D., Chai, B., Caddigan, E., Beck, D. & Li Fei-Fei (2011). Simple line drawings suffice for functional MRI decoding of natural scene categories. Proceedings of the National Academy of Sciences of the USA.

Warren, P. A., Maloney, L. T., & Landy, M. S. (2002). Interpolating sampled contours in 3D: analyses of variability and bias. Vision Research, 42, 2431–2446.

Wertheimer, M. (1923), Untersuchungen zur Lehre von der Gestalt II. Psychologische Forschung, 4, 301-350. Translation published in Ellis, W. (1938). A source book of Gestalt psychology. London: Routledge & Kegan Paul (pp. 71-88).

Wilder, J., Feldman, J. & Singh, M. (2011). Superordinate shape classification using natural shape statistics. Cognition, 119:325–340.

Wilder, J., Singh, M. & Feldman, J. (2013). Detecting shapes in noise: the role of contour-based and region-based representations. Annual Meeting of the Vision Sciences Society (VSS 2013).

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.

Yuille, A.L., Fang, F., Schrater, P., & Kersten, D. (2004). Human and ideal observers for detecting image curves. In S. Thrun, L. Saul, & B. Schoelkopf (Eds.). Advances in neural information processing systems, Volume 16, 59–70. Cambridge, MA: MIT Press.