Technical Report #73, Rutgers University Center for Cognitive Science

# Information along contours and object boundaries

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

Attneave (1954, *Psychological Review*) famously suggested that information along visual contours is concentrated in regions of high magnitude of curvature, rather than being distributed uniformly along the contour. Here we give a novel formal derivation of this claim, yielding an exact expression for information, in Shannon's sense, as a function of contour curvature. Moreover, we extend Attneave's claim to incorporate the role of *sign* of curvature, not just *magnitude* of curvature. In particular, we show that for closed contours, such as object boundaries, segments of negative curvature (that is, concave segments) literally carry greater information than corresponding regions of positive curvature (i.e., convex segments). The psychological validity of our informational analysis is supported by a host of empirical findings demonstrating the asymmetric way in which the visual system treats regions of positive and negative curvature.

In 1954, Attneave proposed that information along a visual contour is not distributed uniformly, but rather is concentrated in regions of high magnitude of curvature.[1] His observation was informal, but astute, and helped to inspire interest in information-processing approaches to the study of vision. Fig. 1a shows a shape with points of locally maximal magnitude of curvature marked. By way of demonstration that such points convey most of the psychologically important information about shape, Attneave drew a line drawing of a cat by taking only the points of local maxima of curvature magnitude, and joining them with straight line segments.[2] The resulting line drawing (now popularly known as 'Attneave's cat') was easily recognizable, suggesting that not much loss of information had occurred. Attneave (1954) also briefly described the results of an experiment in which participants were asked to approximate two-dimensional shapes with a fixed number of points, and then asked to indicate where on the original shapes these points were located. Histogram plots of the points selected revealed salient peaks at precisely the points of local maxima of curvature magnitude (similar to Fig. 1b). The de-

Supported in part by NSF SBR-9875175 to J.F. and NSF BCS-0216944 to M.S. We are very grateful to Howard Resnikoff for thoughtful comments on an earlier draft. Please direct correspondence to Jacob Feldman, Dept. of Psychology, Center for Cognitive Science, Rutgers University – New Brunswick, 152 Frelinghuysen Rd., Piscataway, New Jersey, 08854, or by e-mail at jacob@ruccs.rutgers.edu.

[1] Curvature is sometimes treated as an *unsigned* quantity—the magnitude of the tangent deriviative or the "degree of bendiness"—and sometimes as a *signed* quantity, in which case sign is conventionally assigned positive for turns towards the interior of the "figure" (i.e. convexities) and negative for turns away from the interior (concavities). This discrepant senses can cause confusion, for example when a reference to "low curvature" can refer either to a relatively straight curve (when curvature is used in the unsigned sense) or a region with high magnitude in the negative direction (i.e. a sharp concavity). Attneave used the term curvature in its unsigned sense. Thus in modern language his claim was that information depends on the *magnitude* of curvature. He made no reference in his paper to the sign of curvature, and his proposal did not distinguish between convex and concave regions of a contour.

[2] Irving Biederman (speaking informally at the Psychonomic Society conference, November, 2000) has related that, in Attneave's own telling—and contrary to myth—Attneave never actually made a smoothly curved line drawing of a cat. Rather, Attneave drew the famous feline polygon by hand directly from visual inspection of his own pet.

tails of Attneave's experiment were apparently never published (Attneave's 1954 article cites only a 'mimeographed note'). However Norman, Phillips, and Ross (2001) have recently conducted this experiment and replicated Attneave's results. Moreover, contour deletion experiments (Biederman and Blickle, discussed in Biederman, 1987) have shown that deletion of high-curvature contour segments creates greater difficulties in recognition than deletion of low-curvature segments of comparable length, demonstrating the special role high-curvature contour segments play in recognition.

More recently, Resnikoff (1985) has provided a derivation of Attneave's claim, based on Shannon's mathematical definition of information. Although Resnikoff deserves credit for placing Attneave's proposal on a formal footing for the first time, we feel that his derivation has several problems that leave it short of providing a mathematical substantiation of Attneave's idea (see Appendix). In this article, we provide a novel derivation of the information content of contours, which does not require the assumptions implicit in Resnikoff's analysis, but rather is informed by recent psychophysical findings about the mental representation of curves. Moreover, we extend the informational analysis to the case of *closed contours*—as might correspond to object boundaries—deriving an asymmetry in the information content of negative and positive curvature regions. This analysis extends Attneave's original claim—which treats positive and negative curvature regions symmetrically—and is supported by a host of empirical findings in the literature demonstrating the influence of sign of curvature on shape perception.

## Information

We begin with a statement of Shannon's formula for a continuous measure $M$. Assume first a distribution (probability density function) $p(M)$, which represents the observer's beliefs about the value of $M$ before a measurement is taken. What information is gained by measuring $M$? Shannon's insight was that this depends on the value obtained, and, more specifically, on its likelihood. If the observed $M$ is relatively close to what was expected—say, it was the most likely case—then relatively little information has been gained by measuring it. But if it reveals a surprising value—say, something in the tails of the distribution $p(M)$—then relatively much information has been gained. Specifically, Shannon showed that this dependence must follow the negative logorithm[3] of the probability, i.e.,

$$u(M) = -\log[p(M)]. \tag{1}$$

The quantity $u(M)$ is sometimes called the *surprisal* of $M$. The *information* contained in the distribution $p(M)$, i.e. the entire ensemble of probabilities $p(M)$ taken as a whole, is simply the expected value of the surprisal,

$$I(p) = -\sum_M p(M)\log[p(M)]. \tag{2}$$

that is, the mean of the all possible surprisals weighted by their probabilities.

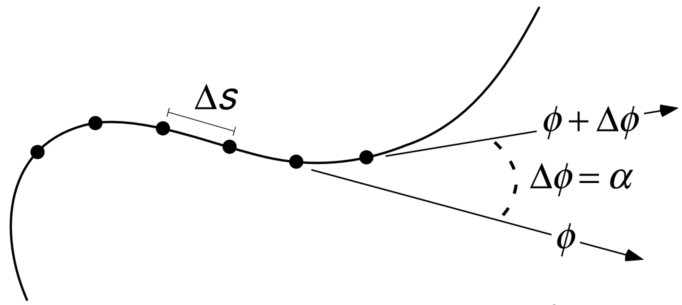


*Figure 2.* A simple plane curve sampled at intervals of arclength $\Delta s$. Each point has a tangent $\phi$; the angle $\Delta\phi$ between successive tangents is denoted $\alpha$.

## Contours

Now consider the case of a simple planar curve (i.e., with no self-intersections) of length $L$, sampled at $n$ uniformly-spaced points separated by intervals $\Delta s = L/n$ (Fig. 2). From point to point along the sampled curve, the tangent direction changes by an angle $\Delta\phi$. (Without loss of generality, we assign the field of normals such that positive values of $\Delta\phi$ correspond to clockwise turns, and negative values to counter-clockwise turns.) The structure of the (sampled) curve is determined completely by the succession of choices of $\Delta\phi$. Let $\alpha$ denote the change in angle $\Delta\phi$ (see Fig. 2). What is the probability distribution of $\alpha$? That is, as one moves around the curve, choosing successively the next change in tangent angle, from what distribution are these choices drawn?

Recent work modeling the process by which human observers interpolate smooth contours through fields of dots suggests an answer to this (Feldman, 1997, 2001; see also Williams & Jacobs, 1997). Experiments suggest that the human visual system assumes a distribution for $\alpha$ that is approximately normal[4] (Gaussian) centered at "straight" ($\alpha = 0$), i.e.

$$p(\alpha) = \mathcal{N}(0,\sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{\left(-\frac{\alpha^2}{2\sigma^2}\right)}, \tag{3}$$

where $\sigma$ is the standard deviation of the distribution, empirically assigned a value of about $\pi/3$ radians (60°). The details of this distribution are actually not very important to

[3] Treatments of information theory usually assume logs in base 2, but the choice of base does not really matter since they differ only by a multiplicative constant. In what follows we actually use base $e$ for reasons that will become apparent.

[4] This equivalence must be approximate because a Gaussian has infinite support, while the required distribution has support $(-\pi,\pi)$. Still the functional form is simple and the approximation is good near $\alpha = 0$, which is the region of psychological interest. It should be clear that while empirical data support the choice of a Gaussian, they likewise support any number of generally similarly-shaped distributions. We adopt the Gaussian because it has approximately the desired functional form and is simple to work with; beyond these factors we do not attach much importance to its particular properties. Elsewhere in the paper we consider the case of arbitrary distributions more explicitly in order to substantiate the claim that our main conclusions do not depend on the Gaussian assumption.

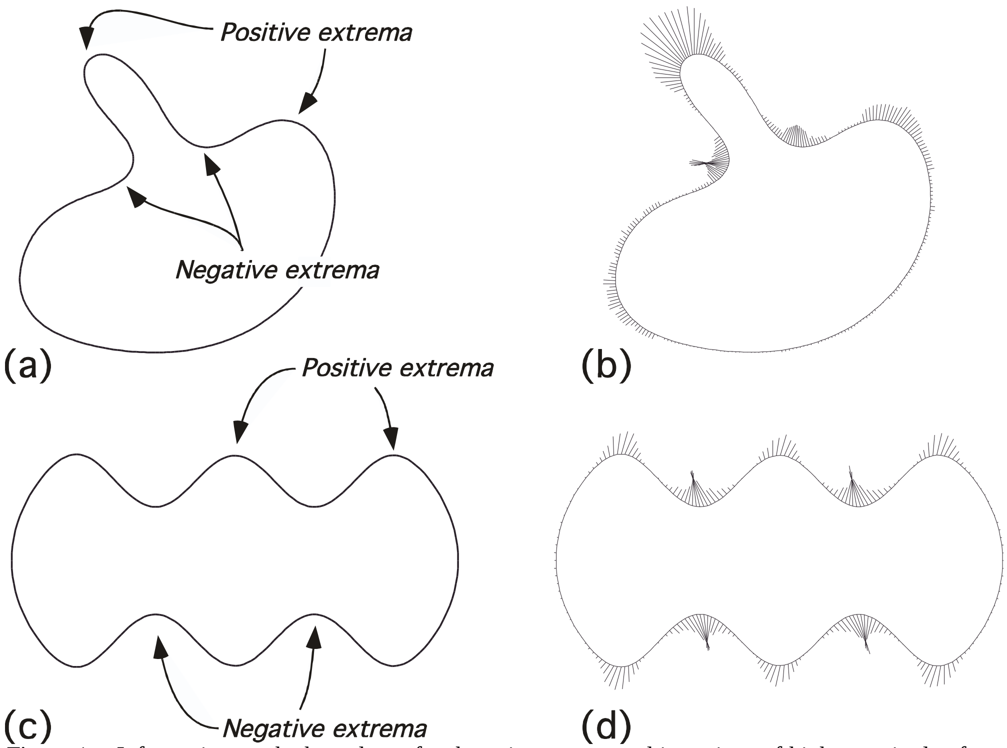


*Figure 1.* Information on the boundary of a shape is concentrated in regions of high magnitude of curvature. (a) A shape with curvature extrema marked, including both positive (convex) extrema and negative (concave) extrema (i.e., minima of signed curvature). (b) The same shape with contour information (surprisal) plotted, reminiscent of Attneave's (1954) histograms. (c) Another shape, with matching regions of positive and negative curvature with equal magnitude, and (d) a plot of surprisal, showing the asymmetry between information due to positive and negative curvature.

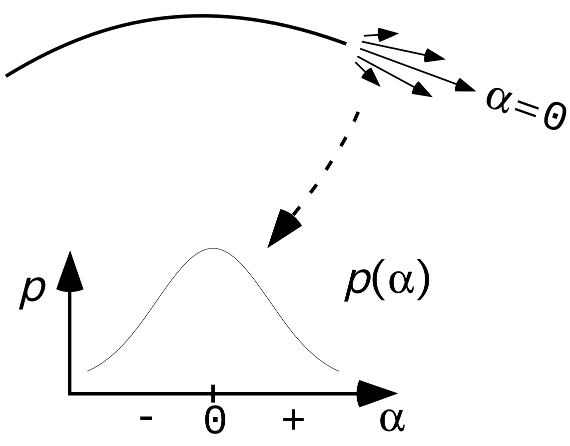


*Figure 3.* The expected change in tangent direction $\alpha$ is normally distributed about 0 (straight).

our argument here (see below); the important properties of this choice of distribution are that (i) it is centered at $\alpha = 0$, meaning that straight continuation of the tangent direction is considered the most likely case, and (ii) probability decreases symmetrically with deviations from straight.

Now, at a particular point along the curve, and particular choice of angle $\alpha$, what is the information at that point? Following Shannon, all we can give for a particular measurement is its surprisal. Combining Eqs. 1 and Eq. 3, we get

$$u(\alpha) = -\log[p(\alpha)] = -\log[\frac{1}{\sigma\sqrt{2\pi}}] + \frac{\alpha^2}{2\sigma^2}. \qquad (4)$$

The first term is an additive constant, not dependent on $\alpha$, which we can ignore. (It gives the absolutely minimal surprisal, obtained in the case of a straight line; its exact value derives from the specifics of the Gaussian prior.) The second term, $\frac{\alpha^2}{2\sigma^2}$, shows how the suprisal depends on $\alpha$: it increases with its square, as measured in standard units of size $\sqrt{2}\sigma$). Fig. 1b shows a plot of the surprisal along a shape boundary, which closely resembles Attneave's empirically-derived histogram plots (see also Norman et al., 2001).

The monotonic increase in suprisal with curvature does not depend on the choice of a Gaussian distribution. To show this, we appeal to Chebyshev's inequality (see Duda, Hart, & Stork, 2001), which provides an upper bound that applies to *all* distributions. One statement of Chebyshev's inequality is that any probability distribution $p$ with mean 0 obeys

$$p(x) \leq \frac{1}{z^2},$$

where $z$ is the $z$-score of $x$, that is, its value normalized by the standard deviation of the distribution. This is a rather loose bound, but one that holds regardless of the details of the distribution. In our notation, it means that for any angular distribution $p(\alpha)$ with mean 0,

$$p(\alpha) \leq \left(\frac{\sigma}{\alpha}\right)^2,$$

where as before $\sigma$ is the standard deviation of the angular distribution. Substituting this bound into the definition of suprisal $u(x) = -\log p(x)$, we see that the surprisal of the turning angle $\alpha$ is bounded below by

$$\begin{aligned} u(\alpha) &\geq -\log\left(\frac{\sigma}{\alpha}\right)^2 \\ &\geq -(\log\sigma^2 - \log[\alpha^2]) \\ &\geq \text{constant} + 2\log|\alpha|. \end{aligned}$$

In words, the surprisal increases with turning angle, at least as quickly as twice the log of its magnitude. Note that the quadratic increase derived above, based on the assumption of Gaussian distribution (see Eq. 8), clearly satisfies this bound. Thus regardless of the exact choice of distribution, information increases monotonically with turning angle.

## *Curvature*

Now we connect this to curvature. The curvature $\kappa$ is the change in tangent direction as we move along the curve, and hence is approximated by the ratio between the change in the tangent direction (i.e., $\alpha = \Delta\phi$) and $\Delta s$:

$$\kappa \approx \frac{\alpha}{\Delta s}. \qquad (5)$$

By definition, this approximation becomes exact in the limit as $\Delta s \to 0$ (i.e., as the number of sample points $n \to \infty$). Note that $\kappa$ inherits its sign from $\alpha$, i.e. clockwise turns are considered positive. Now rearrange terms to yield an expression for $\alpha$:

$$\alpha \approx \Delta s\kappa. \qquad (6)$$

We assumed above that $\alpha$ was distributed normally about 0 with standard deviation $\sigma$ (Eq. 3). Because $\kappa\Delta s \approx \alpha$ this means that $\kappa\Delta s$ is distributed likewise, which in turn means that $\kappa$ is distributed about 0 with standard deviation $\sigma/\Delta s$, i.e.,

$$p(\kappa) \approx \mathcal{N}(0, \sigma/\Delta s) = \frac{\Delta s}{\sigma\sqrt{2\pi}} e^{\left[-\frac{1}{2}\left(\frac{\Delta s\kappa}{\sigma}\right)^2\right]}. \qquad (7)$$

Plugging this into the definition of suprisal (Eq. 1), we find that the surprisal of a given value of curvature $\kappa$ is

$$u(\kappa) \approx -\log\left[\frac{\Delta s}{\sigma\sqrt{2\pi}}\right] + \frac{1}{2}\left(\frac{\Delta s\kappa}{\sigma}\right)^2. \qquad (8)$$

Again ignoring the additive constant (lefthand term), we see that at a given point along a curve the surprisal is proportional to the square of the curvature,

$$u(\kappa) \propto \kappa^2, \qquad (9)$$

and thus increases monotonically with curvature, exactly as Attneave proposed. Moreover, this expression is symmetric with respect to the sign of curvature (i.e., the surprisal is identical for $\kappa$ and $-\kappa$), depending only in its magnitude—again consistent with Attneave's articulation of the claim. The details of Eqs. 8 and 9 depend on the Gaussian assumption, but the Chebyshev argument above can be extended to the curvature case to yield a distribution-free bound on surprisal in terms of curvature,

$$u(\kappa) \geq \text{constant} + 2\log|\Delta s\kappa|. \qquad (10)$$

The main conclusion—that suprisal increases with the magnitude of curvature—is thus guaranteed to obtain regardless of the choice of distribution.

To be slightly more precise, we see that in any of these expressions (e.g. Eq. 8 or 10), information along a contour depends on the product of curvature $\kappa$ and $\Delta s$. What exactly does this mean? Recall that curvature itself is not a scale-invariant quantity. When the entire figure is expanded uniformly by a given ratio (say, by inspecting it from a shorter viewing distance), all curvature values decrease by the same ratio. But because $\Delta s = L/n$, by definition $\Delta s$ scales with the figure. This means that the value $\kappa\Delta s$ is scale-invariant, because whenever the figure doubles in size (say), curvature $\kappa$ is halved but $\Delta s$ is doubled, leaving $\kappa\Delta s$ unchanged. Another way of seeing this is to recall that the magnitude of curvature is equal to the inverse of the radius of the locally best-fitting circle, $1/R$. Hence $\kappa\Delta s = \Delta s/R = L/Rn$. But because $L$ and $R$ scale by the same factor, this ratio is clearly invariant to scale.[5] $\Delta s$ can be thought of as the length of our "measuring stick," and the product $\kappa\Delta s$ as a measure of *scale-invariant curvature* or *normalized cuvature* (see, e.g.,Hoffman & Singh, 1997; Koenderink, 1990).[6]

Hence our expression for the surprisal of curvature (Eq. 8) accords with the intuition that information along a curve is scale-invariant: it depends only on the inherent *shape* of the curve, and not on the particular viewing scale at which we happen to look at it.

[5] Note that this argument does *not* depend on $\Delta s$ being a small or infinitesimal quantity: $L\kappa$ is a measure of scale-invariant curvature for *any* length that is tied to the scale of the figure, as all such measures are clearly proportional to one another.

[6] On 3D surfaces, one has *two* principal curvatures at each point—namely, the curvatures along the directions in which the surface curves the most and the least. Hence, it is possible to define scale-invariant notions of surface curvature by taking ratios of these quantities. Koenderink (1990), for example, defines the *shape index* in terms of the ratio $\frac{\kappa_{max}+\kappa_{min}}{\kappa_{max}-\kappa_{min}}$, a quantity that clearly remains invariant across uniform scalings. For 2D contours, however, each point has a single value of curvature associated with it—and one must thus use some measure of the scale of the figure itself to normalize the value of curvature.

## Closed contours

As we noted earlier, Attneave's claim refers only to the magnitude of curvature, and does not distinguish between positive and negative curvature (i.e., clockwise and counterclockwise turning of the tangent, or equivalently, convex and concave regions). Correspondingly, our Eq. 8 is insensitive to the sign of $\alpha$ or $\kappa$—which followed from the fact that the distribution $p(\alpha)$ is symmetric about 0. So far, there has been no reason for it be otherwise.

However, when a visual contour is the boundary of an object—with one side of the contour assigned "figure" and the other "ground"—an asymmetry is introduced between turning one way and turning the other: one is toward figure (positive curvature), the other toward ground (negative curvature). (Our assumption that clockwise turns have positive sign means that we are travelling clockwise around the figure.) This asymmetry has been demonstrated to have clear psychological consequences.

Citing theoretical analysis and practice from art history, Koenderink (1984) noted that positive curvature regions are typically perceived as having a "thing-like" character, whereas negative curvature regions are perceived as having a "glue-like" character.[7] In their seminal paper on part segmentation, Hoffman and Richards (1984) proposed that the visual system uses negative minima of curvature (points of locally highest curvature magnitude, in concave regions of a shape) to segment shapes into component parts. Thus all curvature maxima (in Attneave's sense of unsigned curvature) are not treated alike psychologically: those with negative curvature are given special status as boundaries between perceived parts, whereas equivalent ones with positive curvature are not (being perceived generally as lying on a single part).

The empirical consequences of this proposed asymmetry between positive and negative curvature (or equivalently, between convex and concave regions) have been demonstrated in a wide variety of tasks, including probe discrimination (Barenholtz & Feldman, in press), positional judgment (Gibson, 1994; Bertamini, 2001), memory for shapes (Driver & Baylis, 1996; Braunstein, Hoffman, & Saidpour, 1989) the perception of figure and ground (Baylis & Driver, 1994; Driver & Baylis, 1996; Hoffman & Singh, 1997), amodal completion (Liu, Jacobs, & Basri, 1999), the perception of transparency (Singh & Hoffman, 1998), and visual search (Hulleman, te Winkel, & Boselie, 2000; Humphreys & Müller, 2000; Elder & Zucker, 1993).

How can the difference between positive and negative curvature be reflected in the informational analysis? Intuitively, the idea is that on a closed contour $C$, with the interior assumed figure, the distribution $p(\alpha)$ is "biased" so that turning in the positive-curvature direction is *more likely* than turning in the negative direction. Otherwise, the curve will not eventually close upon itself. Indeed, the geometry of curves tells exactly *how much* more likely. Over the complete circuit of the curve, the total turning angle must add up to exactly $2\pi$ ($360^\circ$) of total turning angle,

$$\sum_C \alpha = 2\pi, \qquad (11)$$

which means that the expected value (mean) of the distribution $p(\alpha)$, rather than being 0 as before, must now be $2\pi/n$, where $n$ is the number of samples taken at intervals $\Delta s$. For simplicity, we assume the same Gaussian form of the distibution of $\alpha$ as before, except with mean shifted from 0 to $2\pi/n$; that is, the entire distribution is simply translated in $\alpha$-space by a small amount $2\pi/n$ toward the interior of the shape:

$$p(\alpha) = \frac{1}{\sigma\sqrt{2\pi}} e^{\left[-\frac{1}{2}\left(\frac{\alpha-\frac{2\pi}{n}}{\sigma}\right)^2\right]} \qquad (12)$$

Now substituting into the formula for suprisal as before, we get

$$u(\alpha) = -\log\left[\frac{1}{\sigma\sqrt{2\pi}}\right] + \frac{1}{2\sigma^2}\left(\alpha - \frac{2\pi}{n}\right)^2. \qquad (13)$$

Now we progress from angle to curvature by replacing $\alpha$ with its approximation $\kappa\Delta s$, and $\sigma$ with $\sigma/\Delta s$, yielding a formula for the suprisal as a function of curvature:

$$u(\kappa) = -\log\left[\frac{\Delta s}{\sigma\sqrt{2\pi}}\right] + \frac{\Delta s}{2\sigma^2}\left(\kappa\Delta s - \frac{2\pi}{n}\right)^2. \qquad (14)$$

Note that $\kappa$ here must be interpreted in its "signed" sense with positive values assigned to changes of the tangent towards the figure.

Here in the closed-contour case the surprisal is minimal when the tangent direction turns slightly ($2\pi/n$) inwards. Straight ($\kappa = 0$) tangents, rather than being the most expected case as before, are now slightly surprising. The key thing to observe is that points of negative curvature ($\kappa < 0$) are now *more suprising* than points of equivalent positive curvature. However much a given positive value of curvature (i.e., a turn towards the figure) is "in the tails" of the distribution—thus entailing surprise and information—the same value in the negative direction is even *more* in the tails, and hence even more suprising.

This means that negative curvature points literally carry greater information than otherwise equivalent positive-curvature points; Fig. 1d shows a plot of the information (surprisal) along a shape containing convex and concave sections of equal magnitude of curvature, illustrating the asymmetry. The magnitude of contour curvature contributes information, and negative curvature contributes additional information. This picture is supported by recent empirical data showing that perceptual comparisons along the contour are generally slowed by curvature, and slowed an additional amount

[7] Koenderink's analysis was developed in the context 3D-surface curvature—and more precisely, Gaussian curvature—which is somewhat more complicated than contour curvature. However, a theorem by Koenderink (1984) ensures that, for smooth surfaces, the sign of curvature of an occluding contour correspondes to the sign of Gaussian curvature on corresponding surface region. Hence his analysis transfers easily to contour curvature.

by negative curvature, as compared to positive curvature of equal magnitude (Barenholtz & Feldman, in press).

Note that again our main conclusions—that information generally increases with curvature, and is greater for concave compared for convex turns—do not depend on the precise choice of a Gaussian for the distribution (which is, though, supported by empirical data; Feldman, 1997, 2001). Rather they follows directly from the symmetry of the distribution $p(\alpha)$ about its mean, which is required to be positive following the assumption of a closed curve. The same conclusions would have followed from any symmetric monotonically decreasing distribution, although the exact functional form of the resulting equations would be different.

It is especially interesting that no psychological assumptions about the mechanisms underlying part boundary identification were necessary to derive the fact that more information is carried by negative curvature. Rather, this followed simply from the assumption of a closed curve and the implications this must have for the distribution of turning angles as the curve is traversed.

It should be noted that there exists a situation in which the contour is biased to turn *away* from the "figure" rather than toward it: namely, where a simple closed curve bounds a hole or window. In this case, the informational analysis predicts greater concentration of information in regions of the contour that are concave relative to the shaped hole, rather than concave relative to the surrounding material surface. Although this sounds counterintuitive at first, it is actually consistent with recent psychological work on the perception of holes. In particular, holes present the following perceptual anomaly: although the region surrounding the hole is clearly "figural"—in the sense of being a material surface that occludes the backdrop visible through it—the hole is nevertheless seen as a distinct perceptual entity that has its own intrinsic shape (Palmer, 1999). Thus, unlike other forms of "ground," recognition memory for the shapes of holes has been found to be just as good as for similarly-shaped blobs (Rock, Palmer, & Hume, unpublished manuscript; cited in Palmer, 1999, p.286). From the point of view of the visual system, this means that although the surrounding surface is given a figural status as far as depth and occlusion relations are concerned, the hole is given a quasi-figural status, as far as shape analysis is concerned (Nelson & Palmer, 2001; Palmer, 1999; see also Subirana-Vilanova & Richards, 1996). Therefore, it is natural to expect that convexity relationships would be assigned relative to the hole, rather than relative to the surrounding material surface.

## Conclusion

Theories of shape have often emphasized the role of curvature extrema (Richards, Dawson, & Whittington, 1988), and, in the context of perceptual part structure, negative extrema specifically (Hoffman & Richards, 1984; Hoffman & Singh, 1997; Singh, Seyranian, & Hoffman, 1999; Singh & Hoffman, 2001). It follows from our analysis that curvature extrema (in particular, positive maxima and negative minima of signed curvature) are also local maxima of information. Thus in a very concrete sense, these points carry greater information about shape than do other sections of the contour—consistent with Attneave's observation. In addition to providing mathematical justification for Attneave's claim, our analysis also extends it by demonstrating, for closed contours, a role for the sign of curvature. Whereas Attneave considered only the magnitude of curvature—treating regions of positive and negative curvature symmetrically—our analysis shows that regions of negative curvature actually carry greater information than corresponding regions of positive curvature. The psychological validity of this asymmetry is supported by empricial work on the representation of visual shape, which shows that the visual system treats regions of negative and positive curvature quite differently. Finally, our analysis also makes clear that information attaches not to mathematical curvature per se, but rather to a normalized, scale-invariant, version of curvature ($\kappa\Delta s$ in our notation). Thus the contribution of the geometrical structure of a shape to its mental representation does not depend on scale (as curvature proper does); information is a function of "shape only" in the sense of Kendall (1977).

## References


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

Barenholtz, E., & Feldman, J. (in press). Visual comparisons within and between object parts: evidence for a single-part superiority effect. *Vision Research*.

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

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

Biederman, I. (1987). Recognition by components: a theory of human image understanding. *Psychological Review*, *94*, 115–147.

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

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

Duda, R. O., Hart, P. E., & Stork, D. G. (2001). *Pattern classification.* New York: John Wiley and Sons, Inc.

Elder, J., & Zucker, S. (1993). The effect of contour closure on the rapid discrimination of two-dimensional shapes. *Vision Research*, *33*(7), 981–991.

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.

Gibson, B. (1994). Visual attention and objects: Number of objects or convexity? *Journal of Experimental Psychology: Human Perception and Performance*, *19*, 451–470.

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.

Humphreys, G., & Müller, H. (2000). A search asymmetry reversed by figure-ground assignment. *Psychological Science*, *11*(3), 196–201.

Kendall, D. (1977). The diffusion of shape. *Advances in Applied Probability*, *9*, 428–430.

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

Koenderink, J. J. (1990). *Solid shape.* Cambridge: MA: M.I.T. Press.

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

Nelson, R., & Palmer, S. E. (2001). Of holes and wholes: the perception of surrounded regions. *Perception*, *30*, 1213–1226.

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. (1999). *Vision science: photons to phenomenology.* Cambridge: M.I.T. Press.

Resnikoff, H. L. (1985). *The illusion of reality: topics in information science.* New York: Springer-Verlag.

Richards, W., Dawson, B., & Whittington, D. (1988). Encoding contour shape by curvature extrema. In *Natural computation.* Cambridge, MA: M.I.T. Press.

Singh, M., & Hoffman, D. D. (1998). Part boundaries alter the perception of transparency. *Psychological Science*, *9*(5), 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.

Subirana-Vilanova, J. B., & Richards, W. A. (1996). Attentional frames, frame curves and figural boundaries: the inside/outside dilemma. *Vision Reseach*, *36*(10), 1493–1501.

Williams, L., & Jacobs, D. W. (1997). Stochastic completion fields: a neural model of illusory contour shape and salience. *Neural Computation*, *9*, 849–870.

## Appendix: Resnikoff's formulation

Resnikoff (1985) derives an information measure based on contour curvature that, he argues, mimics Attneave's proposal that information is localized in regions of extremal curvature. Resnikoff deserves credit, we feel, for placing Attneave's proposal on a mathematical footing for the first time. However, his derivation has two main problems that leave it short of providing a mathematical substantiation of Attneave's idea. First, his approach is based on the idea of gaining information by making successively finer measurements of a fixed (though unknown) quantity, which seems inappropriate when applied to the problem of measuring contour orientation at successive points along a contour. Second, the behavior of the resulting information measure comes out wrong compared to to both Attneave's claim and other psychological intuitions. In this appendix we briefly review and critique his approach.

Resnikoff's formulation is based on a general framework for quantifying the amount of information gained by successive measurements of a given parameter of fixed, but unknown, value. While Shannon's original theory assumed an observer who knows the underlying probability distribution of messages along the channel (like our shape observer, who we assume to know the distribution of turning angle along the contour), Resnikoff's theory assumes a blank-slate observer lacking this or any other prior information about the quantity in question. The question then is how successive messages (measurements) augment such an observers' knowledge.

Resnikoff's general approach is as follows. Any measurement of a parameter $p$ has finite precision, meaning that it really consists of discovering that the parameter falls within a certain *interval* of finite non-zero size. Assume that a previous measurement has revealed has revealed $p$ to fall within some interval (a,b) of size $|(a,b)|$. Now we take a second measurment and find that $p$ falls within a smaller interval $(a',b')$ of size $|(a',b')| < |(a,b)|$. How much information have we gained by taking the second measurement? Resnikoff shows that the information $I$ (that is, really the surprisal) of the second measurement is

$$I = -\log\left(\frac{|(a',b')|}{|(a,b)|}\right). \tag{15}$$

This expression is very general, showing how information is transmitted via a measurement that increases precision.

Now Resnikoff relates this to curvature by applying Eq. 15 to the measurement of an angle, and specifically, the turning angle $\alpha$ as one moves around a smooth curve at discrete intervals $\Delta s$. Resnikoff considers that as one moves along the curve, successive measurements of the turning angle constitute successive measurements of an angle, suitable for evaluation via Eq. 15. For a given turning angle $\alpha$ and a given reference turning angle $\alpha_R$, this gives

$$I = -\log\left(\frac{\alpha}{\alpha_R}\right), \tag{16}$$

as the information due to a given turning angle $\alpha$ (cf. Resnikoff's Eq. 5.2). Just as in our formulation, this can then be related directly to curvature via the relationship $\alpha = \Delta s \kappa$, to give

$$I = -\log\left(\frac{\kappa}{\kappa_R}\right) \tag{17}$$

as the expression for information as a function of curvature relative to a standard reference curvature $\kappa_R$ (Resnikoff's Eq. 5.8). Resnikoff argues next that, having fixed a standard curvature $\kappa_R$, information will be extremal when curvature is extremal, exactly as Attneave proposed.

However, there are several flaws in the above argument, which we feel make Resnikoff's claim unwarranted. First, application of Eq. 15 to the case of turning angle (or curvature) seems ill-motivated. As derived and developed by Resnikoff, this equation refers to the gain in information by successive measurements of a given *fixed* quantity: that is, to changes in the state of knowledge of the observer about a fixed but unknown parameter. But turning angles at successive points along a curve do not fit this description. Turning angles have different values at different points along the curve because of the inherent geometry of the curve—the fact that it curves at different rates at different points—not because the observer has changed his or her state of knowledge about some fixed quantity. Turning angle decreases (or increases) because the curve bends, not because the observer has measured it more (or less) precisely. Hence applying Eq. 15 to turning angle does not seem valid.

Second, even accepting the validity of Resnikoff's basic set-up, the behavior of his information measure comes out wrong. As Resnikoff notes, his information measure depends always on the comparison (i.e., ratio) of two turning angles (or curvatures). Hence to evaluate the information at a particular point along a curve, one needs first a reference angle to compare it to. There are two general ways of choosing this angle, both of which Resnikoff discusses.

One is to select successive angles as one moves along the curve, comparing each turning angle to the previous one. This leads to information depending not on the turning angle, but rather on the way it (and in the smooth version, the curvature) changes as one moves along the curve. This means, extrapolating to the smooth version, that information would depend on the *derivative* of curvature with respect to arclength—not on curvature itself. This is not what Attneave proposed—and it is not, in fact, psychologically plausible. For example, it would imply that highly curved regions of a contour that were locally nearly circular would contain almost no information.

The second approach, which Resnikoff in any case favors, is to fix a reference turning angle somewhere on the curve and compare all others to it. This way, he argues, information will be extremal when turning angle, and thus curvature, is extremal with respect to this fixed standard. The problem now is that information will be extremal in the wrong way—or more precisely in one of several wrong ways depending on the choice of reference angle. Imagine that we choose a straight (zero-curvature) reference point. Now ratios of other turning angles to the reference will always be infinite (undefined), which is clearly undesirable. So instead, select as a reference a high-curvature point. Now points with similarly-high curvature will have *low* information, while points with low curvature will have high information, exactly the opposite of Attneave's proposal. Finally, consider fixing some low-curvature point as the reference; this is Resnikoff's preference. Now regions of higher curvature will contain more information, with curvature extrema providing the most information, consistent with Attneave's proposal. However straight (zero-curvature) regions will have infinite (undefined) information, which seems qualitatively the wrong behavior.

In our formulation, in contrast with Resnikoff's, the probability of a turning angle derives not from a comparison to another one but by reference to a particular visual expectation about how smooth curves will continue, namely that they will most likely continue straight (in the open-curve case, Eq. 3). Probability is never zero and thus suprisal never infinite.

Indeed, the essential difference between our approach and Resnikoff's concerns the nature of the observer's prior assumptions about the turning angle. In Resnikoff's formulation, all the observer knows when taking a measurement is that a prior measurement revealed it to fall within a particular interval; the observer thus has has no particular expectation about *where* inside that interval the next measurement is likely to fall. This is equivalent to an assumption of *uniform probability density* over the given interval, with all values equally likely. By contrast, in our formulation, we assumed that points had been sampled from a smooth curve, so that probability density about the position of the next point was concentrated in the "forward" direction, at zero turning angle; this assumption was encapsulated in our Gaussian prior. As discussed, this general form (centered at zero and monotonically decreasing away from zero—like a Gaussian though not exclusively so) is supported by empirical data, and, moreoever, is related to the assumption that the points were generated by sampling a smooth curve. Hence in the context of the psychological representation of smooth contours, our non-uniform, forward-centered assumption seems justified.

However, it is well worth noting that in other contexts, something closer to Resnikoff's uniform density assumption might be appropriate. For example, if the series of vertices were generated but by a fractal-like process, with successive angles generated from a uniform density, rather than by sampling from a smooth curve, then Resnikoff's assumptions would be more apt.[8] In this case, information would follow Resnikoff's prescriptions more closely than ours. Of course, the curve resulting from such a process would little resemble the smooth contours discussed above. This raises the fascinating empirical question of whether the human visual system can "tune" its turning-angle distribution to differing environments or contexts, and if so, whether there is any way of empirically measuring the concommitant differences in the information measure. These and other questions await future research.

[8] We are grateful to Howard Resnikoff for this suggestion.