

REPLY

Principles of Contour Information: Reply to Lim and Leek (2012)

Manish Singh and Jacob Feldman
Rutgers University—New Brunswick

Lim and Leek (2012) presented a formalization of information along object contours, which they argued was an alternative to the approach taken in our article (Feldman & Singh, 2005). Here, we summarize the 2 approaches, showing that—*notwithstanding* Lim and Leek’s (2012) critical rhetoric—their approach is substantially identical to ours, except for the technical details of the formalism. Following the logic of our article point by point, Lim and Leek (a) defined probabilistic expectations about the geometry of smooth contours (which they based on differential contour geometry, while we used a discrete approximation—the only essential difference in their approach), (b) assumed that information along the contour was proportional to the negative logarithm of probability, following standard information theory, and then (c) extended this formulation to closed contours. We analyze what they described as errors in our approach, all of which rest on mathematical misunderstandings or bizarre misreadings of our article. We also show that their extension to 3-dimensional surfaces and their “modified minima rule” contain fatal deficiencies.

Keywords: contour, shape, information, curvature, parts

In Feldman and Singh (2005), we presented a formal answer to a question first posed by Attneave (1954): Which portions of a visual contour convey the most information? This problem lies at the foundation of the psychological representation of visual shape, as it depends on the way the brain encodes the geometry of the bounding contours of objects. But aside from an important contribution by Resnikoff (1985), this problem had not received a coherent mathematical answer in the decades since Attneave’s article.¹ In recent years, however, an increasingly sophisticated understanding of the probabilistic representation of visual properties has made an effective approach to this problem possible. In our 2005 article, we proposed a framework for contour information based on a set of very simple ideas, which boil down the following three points.

1. We define probabilistic expectations about the geometry of smooth contours, based on the assumption that smooth contours are most likely to continue to be *straight* at each point, with a decreasing likelihood of larger turning angles (see Figure 1). We expressed this idea mathematically by assuming that the turning angle α follows a von Mises density (similar to a normal or Gaussian density²), $p(\alpha) \propto \exp(\beta \cos \alpha)$. (We also considered

several weaker or more general assumptions; see discussion below.) That is, we assumed that contours tend to continue to be straight with some random deviation about this direction.

2. Following the principles of standard information theory, we assume that the information at each contour point corresponds to the *surprisal* of its turning angle α , defined as $-\log p(\alpha)$. In other words, the shape information at each contour point depends on how “surprising” the geometry at the point is, given our probabilistic expectations $p(\alpha)$ about how the contour will continue. This leads directly to Attneave’s (1954) claim that information content increases with the curvature of the bounding contour.

3. Information along *closed* contours (e.g., object boundaries) can be understood by modifying the distribution of the turning angles $p(\alpha)$ based on the constraint that the contour must eventually close. This yields the mathematical result that points of negative curvature (concavities) convey more information than do points of positive curvature (convexities) because they are more “surprising.” This extends Attneave’s (1954) original claim and is consistent with the psychological finding that convexities and concavities play different roles in shape perception (e.g., Barenholtz & Feldman, 2003; Cohen & Singh, 2007; de Winter & Wagemans, 2008a, 2008b; Hoffman & Richards, 1984; Hoffman & Singh, 1997; Koenderink & Van Doorn, 1982; Panis, de Winter, Vandekerckhove, & Wagemans, 2008) and that changes to shapes

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

The research discussed here was supported by National Science Foundation Grant CCF-0541185; National Institute of Health, National Eye Institute Grant 15888; and National Institute of Health, National Eye Institute Grant 0214924.

Correspondence concerning this article should be addressed to Manish Singh, Department of Psychology and Center for Cognitive Science, Rutgers University—New Brunswick, 152 Frelinghuysen Road, Piscataway, NJ 08854. E-mail: manish@rucss.rutgers.edu

¹ See the Appendix of our original article for a critique of Resnikoff’s (1985) approach.

² The von Mises is the counterpart of the Gaussian or normal density suitable for angular measurements (Mardia & Jupp, 2000). β is a parameter acting like the reciprocal of the variance of a normal distribution. The two functions have provably analogous properties and, in any case, are nearly identical over the range under consideration.

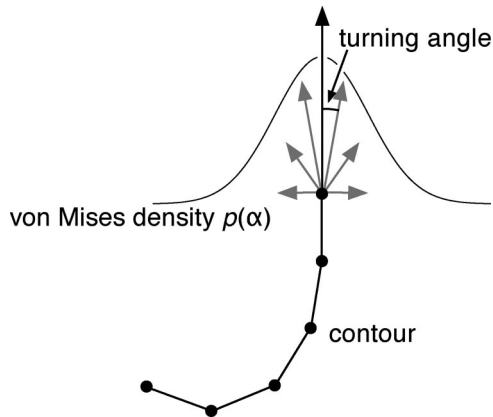


Figure 1. In the Feldman and Singh (2005) approach, a contour is assumed to continue with a turning angle that is distributed as a von Mises (approximately normal) distribution with a mean of 0° (i.e., straight). The information at a point depends on the surprisal of the turning angle, $-\log p(\alpha)$.

are easier to detect at concavities than at convexities (e.g., Barenholtz, Cohen, Feldman, & Singh, 2003; Bertamini & Farrant, 2005; Cohen, Barenholtz, Singh, & Feldman, 2005; Vandekerckhove, Panis, & Wagemans, 2008).

This approach is mathematically coherent and consistent with standard information theory, and moreover, it can easily be extended in several ways. One that we mentioned in the original article is that information along an extended contour C can be understood by integrating surprisal along that contour, yielding a natural measure of contour complexity $-\sum_C \log p(\alpha)$. Another is that information can be extended to include the contribution of higher derivatives of the contour—for example, changes in curvature—merely by placing analogous probabilistic assumptions on those derivatives, such as the assumption that the change in the turning angle (like the turning angle itself) is most likely to be zero. This corresponds to an expectation of cocircularity, that is, a tendency to continue contour curvature, which is supported by empirical evidence in the context of contour integration (e.g., Feldman, 1997; Pizlo, Salach-Goyska, & Rosenfeld, 1997) and the visual extrapolation of contours behind occluders (Singh & Fulvio, 2005, 2007).

Lim and Leek's (2012) Approach

Lim and Leek (2012) presented what they describe as an alternative approach to formalizing contour information, and identified what they describe as errors in our development of the formalism. Notwithstanding their critical rhetoric, their approach echoed our framework 1–3, point by point. The only substantive difference, detailed below, was that their formulation started from probabilistic assumptions about the curvature κ of a smooth contour (a differential property, defined in the limit as arc length goes to zero), while ours began with probabilistic assumptions about the turning angle (a discrete quantity, measurable in the image). Specifically, they assumed a probability density function over curvature, $p(\kappa)$ —analogous to our density over turning angle $p(\alpha)$. From there on, their development proceeded in a manner exactly parallel to ours:

- 1'. Assume contour curvature has distribution $p(\kappa)$ that peaks at zero curvature (i.e., “straight”) and decreases with increasing curvature magnitude.
- 2'. Assume that information depends on the corresponding surprisal, $-\log p(\kappa)$.
- 3'. Show that the above formulation can be modified for the case of closed contours by modifying the distribution $p(\kappa)$ to reflect the constraint that the contour must eventually close.

Lim and Leek (2012) then extended this formulation to three dimensions (see below). Regardless of arguments about the benefits of their continuous formulation versus our discrete one, it is important to recognize that beyond this technical distinction (curvature vs. turning angle), their approach was essentially conceptually identical to ours—hardly the “novel approach” their article promises.

Discrete Versus Continuous Foundations

Contour curvature is a differential property of smooth curves, technically the derivative of the tangent direction with respect to arc length (and therefore determined by the second derivative of the curve). As such, it is defined as a limiting value as one considers the abstract smooth curve in a neighborhood of diminishing size—in effect, it quantifies how rapidly the curve is bending in an infinitesimal neighborhood. By contrast, the turning angle α measures how much the curve has actually “turned” in a given nonzero interval. Basing a theory of shape information on a differential property may be appealing to a mathematician concerned primarily with differential geometry, but we choose to base our approach on discrete quantities that are actually measurable in the image because this ties our work to problems faced by the biological visual system, which can only measure contours with finite resolution. This also allowed us to connect our analysis to the extensive psychophysical and physiological literature on contour integration, which has been an active area of research at least since the early 1970s (e.g., Elder & Goldberg, 2002; Field, Hayes, & Hess, 1993; Geisler, Perry, Super, & Gallogly, 2001; Geisler & Perry, 2009; Lamote & Wagemans, 1999; Smits & Vos, 1987; Uttal, 1973).

In the literature on contour integration, researchers have examined how the visual system groups discrete local elements (e.g., dots, oriented edges, or Gabor patches) into representations of extended contours. Among other things, this literature has documented the geometric relations—most importantly, the turning angles between neighboring discrete elements—that support the percept of an extended contour. (Indeed, by definition, one cannot actually study empirical properties of contour curvature without going via a discrete approximation such as turning angle.) These psychophysical findings have led to the important idea of an “association field” (Field et al., 1993; Geisler et al., 2001), have been linked to the statistics of natural images (Elder & Goldberg, 2002; Geisler & Perry, 2009; Geisler et al., 2001), and have informed physiological findings about long-range horizontal connections between orientation-sensitive units in the primary visual cortex (e.g., Bosking, Zhang, & Fitzpatrick, 1997). Performance

on contour integration tasks has also allowed researchers to estimate the distribution of the turning angles that the visual system implicitly expects. Specifically, the Gaussian or von Mises distribution of the turning angle has been documented in a number of studies (Feldman, 1996, 1997, 2001). Hence, our assumption of such a distribution, unlike Lim and Leek's (2012) adoption of a distribution over curvature itself, was not simply a mathematical abstraction but a scientifically documented premise.

Moreover, as Lim and Leek (2012) eventually acknowledged, unlike turning angle, curvature is not scale invariant. This means that any information theoretic treatment would, likewise, have information varying with the scale, or overall size, of the object. Lim and Leek only mentioned this late in their development and asserted that it can be patched, but their solution is ad hoc—repairing a problem that stems from the foundation of the theory—and without empirical support.

Technical Discrepancies

In previous articles on turning angle, Feldman (1996, 1997, 2001) found psychophysical evidence for a Gaussian or von Mises density for the turning angle α , which is corroborated by physiological work on the tuning properties of orientation-selective neurons in the primary visual cortex (Swindale, 1998). But in Feldman and Singh (2005), we noted that certain key claims about contour information do not depend on the specific form of the distribution but, instead, were guaranteed as long as the distribution $p(\alpha)$ peaked at zero (in the open-contour case) and decreased monotonically with increasing turning angle. Indeed, as we stated in the original article, the fact that contour information rises with turning angle stems directly from the assumption of decreasing probability with larger turning angles, via the very basic relation that surprisal is simply $-\log p(\alpha)$, an assumption that Lim and Leek (2012) explicitly adopted. Indeed this basic observation—that information along contours depends on probabilistic expectations about how curves continue—was the principal contribution of our article, and this seems to have been embraced wholeheartedly by Lim and Leek (2012).

Lim and Leek (2012) did dispute our more specific claim that there is a bound on information as a function of curvature regardless of distributional form, which stems from a much more technical argument. Our claim was that the probability of more extreme turning angles is bound by a limit that diminishes with the increasing turning angle because of what is called the *Chebyshev bound*. The Chebyshev bound shows that the tail area of any probability distribution must diminish as one goes further out in the tails, regardless of the form of the distribution. What this means in our context is that the probabilities of turning angles fall under a bound that decreases monotonically as they get more extreme. Lim and Leek (2012) described our derivation as mathematically flawed, but their argument stems from a puzzling confusion between the terms “probability” and “probability density.” Probability density is a quantity associated with a continuous random variable, such as the turning angle, which must be integrated over some nonzero interval in order to yield a nonzero probability. Unlike probability, probability density can take values higher than 1, so long as the integral over any interval—the probability itself—does not exceed 1. For smooth densities over continuous random variables (like the one over turning angle

considered in our article), the “probability” per se of a single point value—say, 45° —is automatically 0 (because, again, it must be integrated across some nonzero interval to amount to a nonzero value). Hence, an expression such as “the probability of an angle α_0 ” really means “the integral of the probability density function in a small neighborhood of α_0 ,” more properly notated $\int_{\alpha_0-\epsilon}^{\alpha_0+\epsilon} f(\alpha) d\alpha$, where $f(\cdot)$ is the probability density function—a notation we avoided in our article as pedantic and potentially confusing. In words, what this means is that the probability of a given angle corresponds to the area under the density curve over a small range (e.g., the probability that the turning angle falls between 45° and 46°)—an area that is a *subset* of the tail and thus necessarily smaller than the tail, which Chebyshev showed is itself bounded in size. This was our original argument, and it is correct. Lim and Leek's suggestion that it leads to an absurd conclusion—that there is no bound on the magnitude of probability—depends entirely on perversely interpreting $p(\alpha)$ as *density* rather than probability, which results only in showing that probability *density* can be larger than 1—which is correct, but trivial. Once one distinguishes correctly between *probability* and *probability density*, Lim and Leek's Equation 3, which stated that the area under the curve in the neighborhood near a given angle, $p(t + \Delta t)\Delta t$ in their notation, is limited by the Chebyshev bound, becomes *precisely* equivalent to our assertion that they claim to dispute. In sum, the only way our claim could appear to be invalid is if one misreads the notation $p(\alpha)$ to mean *probability density at an infinitesimal point*—an interpretation that is no way supported by our article and is very unlikely to be drawn by any mathematically knowledgeable reader. The authors' point is therefore most charitably understood as a pedantic criticism about notation and less charitably understood as a deliberate obfuscation intended to create the appearance of an error.

Returning to the main issue, the primary contribution of our original article was the observation that information along contours can be understood as Shannon surprisal ($-\log p$), once one adopts a suitable probability model for contours. One can make a range of assumptions about this probability model, with concomitant constraints on the resulting surprisal. We can divide these assumptions into three basic cases, all of which were mentioned at various points in our original article, but which Lim and Leek (2012) conflated in their discussion:

- Strongest Assumption

One can assume a particular distribution of turning angle, such as that it is Gaussian or von Mises. As documented above and in our original article, such an assumption is well supported by the psychological literature. In this case, surprisal clearly increases with turning angle. Lim and Leek (2011) did not dispute this.

- Weaker Assumption

Instead of assuming a particular functional form for the distribution, we can simply assume that the distribution has a probability that is monotonically decreasing with the increasing turning angle. In this case, the increase in surprisal $-\log p$ with increasing turning angles again follows immediately.

- Weakest Assumption

Finally, instead of assuming a monotonically decreasing probability distribution, we can instead make *no* assumptions about the functional form of the distribution, in which case (the Chebyshev argument shows) one can still make a surprisingly strong claim about angle probability, namely that it is *bound* by a monotonically decreasing function (and hence that contour information is bound by a monotonically increasing function). Lim and Leek (2012) disputed this argument, but in so doing, they repeatedly conflated it with one of the above assumptions. (They also conflate “probability” and “probability density,” as we noted above.) The fact that the bound decreases monotonically obviously does not mean that the function itself decreases (which is why our original article asserted only that the *bound* decreases in this case). Of course, in practice, a decreasing bound places a strong constraint on likely angles since it implies that the range of possible angle probabilities gets tighter as the angle gets larger.

Another point of criticism raised by Lim and Leek (2012) is that as the sampling density of the contour is increased, the effect of the turning angle on information content gets progressively smaller. This criticism is equally misplaced. For any finite-resolution measurement system, it must be the case that the smaller the window through which one examines a smooth contour, the closer to linear it looks. If this microscopic view of the contour is all that is available to the system, it would indeed have to conclude that that contour segment is (locally) flatter. Their criticism thus simply boils down to the fact that in a sufficiently small neighborhood, all smooth curves approximate straight lines. From the point of view of perception, this is uninteresting because the visual system does not have access to an arbitrarily fine view.³ Moreover, their criticism presupposes that when the separation between consecutive sample points is decreased, the dispersion (or spread) of the von Mises distribution on turning angle will remain unchanged. This assumption is unjustified. Simple considerations of the geometry of smooth contours suggest that the distribution of turning angles would be expected to be tighter (smaller spread) when the contour is examined through smaller apertures. This criticism of Lim and Leek’s (2012) is thus based on a straw man—a patently invalid extension of our framework, which serves only to confuse. Finally, Lim and Leek’s complaint about our visualization code is trivial—analogueous to complaining about the scale of the *y*-axis in a data plot.

Lim and Leek’s (2012) Three-Dimensional Extension and the “Modified Minima Rule”

Later in their article, Lim and Leek (2012) attempted to generalize their approach from two-dimensional contours to three-dimensional (3D) surfaces. Unfortunately, their proposed 3D extension fails several basic tests of plausibility.

Curvature on a 3D surface is substantially more complex than for a contour, and several alternative formulations of curvature are possible (Koenderink, 1990). At any given point on a surface, the curvature of a path depends on the direction in which it is taken; there is a different curvature value for each possible direction. Of these directions, one has maximal curvature and one has minimal curvature, and these form an orthogonal basis for describing surface curvature. These two directions determine the so-called principal curvatures, κ_1 and κ_2 , which must be combined in some way

in order to characterize “the curvature” of the surface at that point. Different measures of surface curvature combine them in different ways: *Gaussian curvature* takes their product and *mean curvature* takes their average, whereas *curvedness* and *shape index* use more complex combination functions (Koenderink, 1990). A central challenge in defining information content for surfaces is to determine how exactly the two principal curvatures contribute to the surprisal. Lim and Leek (2012) unfortunately disregarded this deep issue altogether and simply adopted Gaussian curvature as a basis for their information measure—giving no motivation or principled reason for their choice and indeed not even mentioning that alternative measures of surface curvature exist.

Unfortunately, Gaussian curvature is not in our view an adequate basis upon which to quantify information. Most obviously, Gaussian curvature is zero if *either* of the two principal curvatures is zero—such as *everywhere* along the curved surface of a cylinder. More generally, Gaussian curvature is zero everywhere along the surface of a *generalized cylinder* with a straight axis but an arbitrary (though fixed) cross-section—regardless of the curvature profile of the cross-section (e.g., the cookie-cutter cat in Figure 2A). This means that Lim and Leek’s (2012) formalism predicts that all points on such a curved surface convey identical—and minimal—shape information. Similarly, the surface in Figure 2B, although it has ridges and valleys of varying strengths, nevertheless has zero Gaussian curvature everywhere (because one of the two principal curvatures is consistently zero). That these variations in surface geometry convey *no* variation in shape information blatantly contradicts intuition, and this is not supported by any empirical finding of which we are aware. It is disappointing that these rather obvious defects in the formulation are not acknowledged, much less addressed, in Lim and Leek’s article.⁴

Lim and Leek (2012) went on to argue that their measure of 3D surface information *implies* a constraint on 3D part boundaries, namely, that they can only occur within regions of negative Gaussian curvature (saddle-shaped regions). Indeed, the tendency for such regions to contain part boundaries was observed almost three decades ago in a well-known article by Koenderink and van Doorn (1982), which Lim and Leek did not cite. Lim and Leek gave no argument whatsoever for any mathematical connection between shape information as they have quantified it and part boundary status—which makes it difficult to evaluate its intended scope.

In any case, it is readily apparent that part boundaries need not fall *only* within negative Gaussian curvature regions, as Lim and

³ A more elaborate treatment of the information content of contours should of course take into account analysis at multiple scales. Although the details of such an analysis await future work, our original article set up the mathematical components that would be required for it.

⁴ It may be argued that surfaces with zero Gaussian curvature are nongeneric because one of the principal curvatures must be exactly zero. However, whether such surfaces are generic depends on one’s generative model for surfaces. If the distribution on each principal curvature peaks sharply at 0, for example, such surfaces will typically *not* be nongeneric. Further, small perturbations to the zero principal curvature on such surfaces will perturb the Gaussian curvature only slightly away from zero. Hence, even taking small perturbations into account, it would still be the case that variations in the nonzero principal curvature (such as the outline of the cookie-cutter cat in Figure 2A) will affect the surface surprisal only minimally.




Figure 2. Examples of 3D shapes with zero Gaussian curvature everywhere and thus—in Lim and Leek’s (2012) account—constant, minimal, shape information everywhere. In both shapes, Gaussian curvature is zero everywhere because one of the two principal curvatures is zero. Under Lim and Leek’s account, the cookie-cutter cat in A, unlike Attneave’s (1954) cat, would have *no* variation in the concentration of shape information; similarly, in B, the variations in the strengths of the ridges and valleys would convey no variation in shape information. These predictions blatantly contradict intuition.

Leek’s (2012) “modified minima rule” asserts. Many biological limbs have boundaries that do *not* fall exclusively within negative Gaussian curvature regions but are nevertheless perceived as psychological parts of objects (see, e.g., de Winter & Wagemans, 2006; Singh, Seyranian, & Hoffman, 1999). An example is the human shoulder joint, which has negative Gaussian curvature underneath but positive Gaussian curvature above. Thus, Lim and Leek’s definition of *part* would seem to exclude human arms, as well as legs and most animal limbs. Indeed, it is well established that the loci of negative minima are by themselves insufficient to segment shapes into parts, and various researchers have proposed geometric factors beyond negative minima for part decomposition (e.g., de Winter & Wagemans, 2006; Siddiqi, Tresness, & Kimia, 1996; Singh & Hoffman, 2001; Singh et al., 1999). By contrast, Lim and Leek’s proposal simply added a further constraint on *which* negative minima should be used (namely, those that lie in hyperbolic regions). It does nothing to address the limitations of negative minima that have been discussed in the parts literature, and thus, it contributes little to our understanding of part decomposition.

The surfaces in Figure 2 illustrate another shortcoming of limiting part boundaries to negative Gaussian curvature regions: These surfaces contain no negative Gaussian curvature regions at all but do contain several fairly salient part boundaries (e.g., at the base of the ears or tail of the cookie-cutter cat and along the valleys of “wave”).

Conclusion

In summary, we are gratified that Lim and Leek (2012) seem to agree with the major elements of the approach we proposed in our 2005 article, though we are somewhat baffled by the negative way in which they have framed their agreement. Our approach is simple, is internally coherent, agrees with both intuition as well as psychophysical and physiological data on contour perception, and harmonizes with the growing literature on probabilistic approaches to visual perception. The defects that Lim and Leek listed stem either from mathematical misunderstandings or from bizarre mis-

readings of our article—some of which we hope we have clarified above. Unfortunately, while the prospect of extending this approach to 3D is appealing, the simplistic manner they propose to do so, based solely on Gaussian curvature, is unconvincing. Still, we welcome attention to this important problem and look forward to what we hope will be more edifying developments in the future.

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*, 1–9.
- Barenholtz, E., & Feldman, J. (2003). Visual comparisons within and between object parts: Evidence for a single-part superiority effect. *Vision Research*, *43*, 1655–1666.
- Bertamini, M., & Farrant, T. (2005). Detection of change in shape and its relation to part structure. *Acta Psychologica*, *120*, 35–54.
- Bosking, W., Zhang, Y. B. S., & Fitzpatrick, D. (1997). Orientation selectivity and the arrangement of horizontal connections in the tree shrew striate cortex. *The Journal of Neuroscience*, *17*, 2112–2127.
- Cohen, E. H., Barenholtz, E., Singh, M., & Feldman, J. (2005). What change detection tells us about the visual representation of shape. *Journal of Vision*, *5*, 313–321.
- Cohen, E. H., & Singh, M. (2007). Geometric determinants of shape segmentation: Tests using segment identification. *Vision Research*, *47*, 2825–2840.
- 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*, 50–64.
- Elder, J. H., & Goldberg, R. M. (2002). Ecological statistics of Gestalt laws for the perceptual organization of contours. *Journal of Vision*, *2*, 324–353.
- Feldman, J. (1996). Regularity vs. genericity in the perception of collinearity. *Perception*, *25*, 335–342.
- Feldman, J. (1997). Curvilinearity, covariance, and regularity in perceptual groups. *Vision Research*, *37*, 2835–2848.
- Feldman, J. (2001). Bayesian contour integration. *Perception & Psychophysics*, *63*, 1171–1182.
- Feldman, J., & Singh, M. (2005). Information along contours and object boundaries. *Psychological Review*, *112*, 243–252.
- 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*, 173–193.
- 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). Saliency of visual parts. *Cognition*, *63*, 29–78.
- Koenderink, J. J. (1990). *Solid shape*. Cambridge, MA: MIT Press.
- Koenderink, J. J., & van Doorn, A. J. (1982). The shape of smooth objects and the way contours end. *Perception*, *11*, 129–137.
- Lamote, C., & Wagemans, J. (1999). Rapid integration of contour fragments: From simple filling-in to parts-based shape description. *Visual Cognition*, *6*, 345–361.
- Lim, I. S., & Leek, E. C. (2011). Curvature and the visual perception of

- shape: Theory on information along object boundaries and the minima rule revisited. *Psychological Review*. Advance online publication. doi: 10.1037/a0025962
- Mardia, K. V., & Jupp, P. E. (2000). *Directional statistics*. New York, NY: Wiley.
- 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.
- Pizlo, Z., Salach-Goyska, M., & Rosenfeld, A. (1997). Curve detection in a noisy image. *Vision Research*, *37*, 1217–1241.
- Resnikoff, H. L. (1985). *The illusion of reality: Topics in information science*. New York, NY: Springer-Verlag.
- Siddiqi, K., Tresness, K. J., & Kimia, B. B. (1996). Parts of visual form: Psychophysical aspects. *Perception*, *25*, 399–424.
- Singh, M., & Fulvio, J. M. (2005). Visual extrapolation of contour geometry. *Proceedings of the National Academy of Sciences, USA*, *102*, 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. (2001). Part-based representations of visual shape and implications for visual cognition. In T. Shipley & P. Kellman (Eds.), *From fragments to objects: Segmentation and grouping in vision: Advances in psychology* (Vol. 130, pp. 401–459). New York, NY: Elsevier.
- Singh, M., Seyranian, G. D., & Hoffman, D. D. (1999). Parsing silhouettes: The short-cut rule. *Perception & Psychophysics*, *61*, 636–660.
- Smits, J. T., & Vos, P. G. (1987). The perception of continuous curves in dot stimuli. *Perception*, *16*, 121–131.
- Swindale, N. V. (1998). Orientation tuning curves: Empirical description and estimation of parameters. *Biological Cybernetics*, *78*, 45–56.
- Uttal, W. R. (1973). The effect of deviations from linearity on the detection of dotted line patterns. *Vision Research*, *13*, 2155–2163.
- Vandekerckhove, J., Panis, S., & Wagemans, J. (2008). The concavity effect is a compound of local and global effects. *Perception & Psychophysics*, *69*, 1253–1260.

Received September 23, 2011

Revision received April 16, 2012

Accepted April 20, 2012 ■