

Bayesian estimation of the shape skeleton

Jacob Feldman, and Manish Singh

PNAS published online Nov 13, 2006;
doi:10.1073/pnas.0608811103

This information is current as of November 2006.

E-mail Alerts

This article has been cited by other articles:
www.pnas.org#otherarticles

Rights & Permissions

Receive free email alerts when new articles cite this article - sign up in the box at the top right corner of the article or [click here](#).

Reprints

To reproduce this article in part (figures, tables) or in entirety, see:
www.pnas.org/misc/rightperm.shtml

To order reprints, see:
www.pnas.org/misc/reprints.shtml

Notes:

Bayesian estimation of the shape skeleton

Jacob Feldman* and Manish Singh

Department of Psychology, Center for Cognitive Science, Rutgers, The State University of New Jersey, Piscataway, NJ 08854

Communicated by Charles R. Gallistel, Rutgers, The State University of New Jersey, Piscataway, NJ, October 6, 2006 (received for review May 19, 2006)

Skeletal representations of shape have attracted enormous interest ever since their introduction by Blum [Blum H (1973) *J Theor Biol* 38:205–287], because of their potential to provide a compact, but meaningful, shape representation, suitable for both neural modeling and computational applications. But effective computation of the shape skeleton remains a notorious unsolved problem; existing approaches are extremely sensitive to noise and give counterintuitive results with simple shapes. In conventional approaches, the skeleton is defined by a geometric construction and computed by a deterministic procedure. We introduce a Bayesian probabilistic approach, in which a shape is assumed to have “grown” from a skeleton by a stochastic generative process. Bayesian estimation is used to identify the skeleton most likely to have produced the shape, i.e., that best “explains” it, called the maximum *a posteriori* skeleton. Even with natural shapes with substantial contour noise, this approach provides a robust skeletal representation whose branches correspond to the natural parts of the shape.

computation | vision

Skeletal representations of visual shape, in which a shape contour is represented in terms of local symmetries about a set of possibly curving axes, have played a prominent role in theories of visual shape ever since the introduction of the medial axis transform (MAT) by Blum (1) and Blum and Nagel (2). The MAT is widely suspected to play a role in cortical representations of visual shape, perhaps by a neural implementation of Blum’s local “grassfire” procedure. Cells as early as primary visual cortex show enhanced sensitivity near medial points (3, 4), suggesting an early locus of computation. Moreover, medial axes have known psychophysical correlates, including increased sensitivity to contrast (5, 6) and position (7) and probe comparisons in which latency patterns respect perceived axial structure (8). Medial axes are also deeply intertwined in theories of how shapes are decomposed into parts (9); for example, despite considerable controversy about how part cuts (boundaries between perceptually distinct parts) are determined, there is substantial agreement that part cuts must cross a medial axis (10). More broadly, many higher-level theories of shape and shape recognition are substantially based on axial representation of parts (11–13) presupposing prior computation of some sort of skeletal shape representation.

However, the computation of the medial axis skeleton suffers from several notorious problems, including spurious axial branches stemming from hypersensitivity to perturbations along the contour, and counterintuitive results (forking) at the ends of blunt parts (see Fig. 3 *Insets*). More recent advances in the computation of the MAT (14–18) have reduced, but not eliminated, these problems, which seem to be endemic to the underlying geometric conception of the MAT.

Summary of the Approach

The basic idea behind our approach is that real shapes owe their structure to a mixture of generative and random factors, e.g., shapes that are the result of an underlying skeleton plus a stochastic growth process. We apply Bayesian estimation to the problem of identifying a shape’s most likely “generative skeleton,” under simple assumptions about the probability distribution of skeletons (providing a Bayesian prior), and a stochastic

model of how shapes are generated from skeletons (providing a Bayesian likelihood function). The prior favors simple skeletons with relatively few and relatively straight branches. The likelihood model, i.e., the shape-generating stochastic process, assumes that shapes are generated by a lateral outward growth process in which there is some random variation in the direction of growth away from the axis and some random variation in the extent of growth. We then combine this prior and likelihood function by Bayes’ rule, identifying the generative skeleton that is most likely to have produced the shape. An axial branch is included in this skeleton only when the additional skeletal complexity it creates is more than offset by the improved “goodness of fit” to the shape. The estimated skeleton, called the maximum *a posteriori* (MAP) skeleton, is the skeletal interpretation that, under the generative assumptions underlying the prior and likelihood functions, best “explains” the shape.

Bayesian Formulation, Priors, and Likelihood Functions

We begin by assuming a shape given by a discrete approximation $\text{SHAPE} = \{x_1, x_2, \dots, x_N\} \subset R^2$. (We assume a closed shape, but formally all that is required is a boundary with figure and ground assigned, so that the direction of the field of normals is well defined.) Skeletons are generated under a probability density function $p(\text{SKEL})$; and in turn shapes are generated from skeletons under a conditional probability density function playing the role of a likelihood function $p(\text{SHAPE}|\text{SKEL})$. The key idea is that this likelihood function expresses a generative model of shape (19) so that selecting a particular skeletal interpretation, a particular generative skeleton, amounts to explaining the observed shape in the most plausible way under the assumed generative model. Following the Bayesian approach (20–22) the plausibility of a particular skeletal description corresponds to its posterior probability, given by Bayes’ rule:

$$p(\text{SKEL}|\text{SHAPE}) = \frac{p(\text{SHAPE}|\text{SKEL})p(\text{SKEL})}{\sum_i p(\text{SHAPE}|\text{SKEL}_i)p(\text{SKEL}_i)}, \quad [1]$$

summing over all possible skeletons SKEL_i . Because the denominator in this expression is constant for a given shape, we can maximize it by maximizing the numerator, i.e., the product of the prior and likelihood.

Priors. A skeleton SKEL consists of a set of axial segments $C_1 \dots C_K$, hierarchically organized into a root contour, branches, subbranches, etc. We define a prior probability density $p(\text{SKEL})$, using a natural hierarchical extension of our earlier work on contour information (23). For each axial segment C_i , we induce a prior density $p(C_i)$ by assuming that successive points in its discrete approximation are generated by a density function

Author contributions: J.F. and M.S. performed research and wrote the paper.

The authors declare no conflict of interest.

Freely available online through the PNAS open access option.

Abbreviations: MAT, medial axis transform; MAP, maximum *a posteriori*; DL, description length.

*To whom correspondence should be addressed. E-mail: jacob@rucss.rutgers.edu.

© 2006 by The National Academy of Sciences of the USA




Fig. 1. Illustration of the prior probability density over skeletons $p(\text{SKEL})$, illustrating how probability decreases as skeletons branch and/or as axial branches bend.

centered on straight (zero curvature) continuation of the axis, with angular deviation from collinearity α following a von Mises distribution $V(0^\circ, b) \propto e^{\cos(b\alpha)}$ [similar to a Gaussian (normal) distribution but suitable for angular measurements (24)], which has proved accurate in modeling human contour perception (23, 25). Under this assumption, relatively straight axes (α near zero) have high probability, whereas probability decreases with larger turning angles, i.e., with larger magnitude of curvature in the underlying curve. Successive turning angles are assumed independent, so the prior $p(C)$ for a curve C containing a series $\alpha_1, \alpha_2, \dots$ of turning angles in its discrete approximation is $\prod p(\alpha_i)$. To induce a prior over skeletons, we simply augment this prior by assuming that axial branches C_i sprout with fixed probability p_C , which yields the probability of a skeleton SKEL comprising K axes of

$$p(\text{SKEL}) = p_C^K \prod_i p(C_i). \quad [2]$$

This prior is high for skeletons with few and relatively straight axes and diminishes with increased branching or increasing curvature in any of the axial branches (Fig. 1), an assumption




Fig. 2. Likelihood model. (a) Illustration of the likelihood function, $p(\text{SHAPE}|\text{SKEL})$, showing how a shape is generated stochastically from a skeleton. Ribs sprout from each on both sides of each axis, in directions that are perpendicular (normal) to the axis plus a random directional error ϕ , chosen independently for each rib. The ribs have lengths $\mu_C + \epsilon$, where μ_C is the rib length function of axis C , and ϵ is a Gaussian error chosen independently for each rib. The generated shape is the shape formed by the rib endpoints. (b) A more complex shape and its MAP skeleton, showing estimated ribs (i.e., axis-shape correspondences), color-coded to illustrate common axial sources. The distinct colors correspond well to perceptually distinct parts of the shape.

empirically validated by the generally simple forms exhibited by naturally occurring shapes (26). For a skeleton consisting of a single axis, the prior reduces to the established prior for a simple open contour (23) as seems natural.

Likelihoods. The next step in a Bayesian account is the adoption of a likelihood model, in this case meaning a stochastic generative model by which a shape is produced from a hypothesized skeleton. To capture the idea that the shape is “extruded” laterally from the generative skeleton, we assume that from each point on each skeletal axis, “ribs” sprout on both sides, approximately perpendicular to the axis (hence, primarily outward), but with partly random lengths and directions (Fig. 2). More specifically, each rib sprouts in a direction that is perpendicular to the axis plus a random directional error ϕ_x , chosen independently for each rib (i.e., the rib ending at shape point x) from a von Mises density centered on zero, i.e., $\phi_x \sim V(0, b_\phi)$ with spread parameter b_ϕ . The expected rib length μ at each point v along the axial segment C is given by a “rib length function” $\mu_C(v_x)$, which we estimate from the shape assuming only a continuity constraint (see *Methods*). To this expected rib length $\mu_C(v_x)$ is added a random rib length error ϵ_x , chosen independently for each rib from a normal distribution, $\epsilon \sim N(0, \sigma_\epsilon^2)$. The expected rib lengths $\mu_C(v_x)$ are themselves drawn from an exponentially decreasing density $p(\mu) \propto e^{-\beta\mu}$ with decay constant β , meaning that wider axial parts are less likely than narrower ones, with probability decaying gradually with increasing widths. For each shape point x , the expected rib length μ , directional error ϕ , and rib length error ϵ are mutually independent, so the likelihood of the shape point $p(x|\text{SKEL})$ generated by a rib at point v_x along axis C is given by the product

$$p(x|\text{SKEL}) = p(\mu_C(v_x))p(\phi_x)p(\epsilon_x). \quad [3]$$

The likelihood of the entire shape is the product of the likelihoods of its constituent points,

$$P(\text{SHAPE}|\text{SKEL}) = \prod_{x \in \text{SHAPE}} p(x|\text{SKEL}). \quad [4]$$

The MAP Skeleton

Given the prior and likelihood defined as above, the final step is to compute the skeletal structure with maximum posterior probability, the MAP skeleton. We propose estimation of this skeleton as a “competence” or computational theory of mental shape representation, meaning a specification of the function that the human system is attempting to compute when it represents shape (rather than an account of the implementation it uses to compute it). We can maximize the posterior by, equivalently, choosing the skeleton that minimizes the negative logarithm of the posterior, often referred to as its description length (DL) because it reflects the complexity of expressing the hypothesis in an optimal code (27). Taking the negative logarithm of Eq. 1, the DL of the skeletal posterior is just

$$\begin{aligned}
DL(\text{SKEL}|\text{SHAPE}) &= -\log[p(\text{SKEL})] \\
&\quad -\log[p(\text{SHAPE}|\text{SKEL})] + \text{const} \\
&= DL(\text{SKEL}) + DL(\text{SHAPE}|\text{SKEL}) \\
&\quad + \text{const.} \quad [5]
\end{aligned}$$

Apart from the constant term (the negative logarithm of the denominator in Eq. 1) the DL has two additive components: $DL(\text{SKEL})$, which reflects the complexity of the skeletal hypothesis itself, and $DL(\text{SHAPE}|\text{SKEL})$, which reflects the complexity of the shape as described by that skeleton. The MAP skeleton is the description that minimizes the sum of these two complexities. Hence the MAP skeleton is naturally regarded as identifying the simplest description of the shape as the outcome of a skeletal generative process. This attractive interpretation stems directly from the Bayesian conception and is not shared by other stochastic techniques for skeletal-axis computation.

The process of estimating the MAP skeleton requires inverting the likelihood function by choosing, for each shape point, the skeletal point that has “responsibility” for it, i.e., assigns it the highest likelihood. This skeletal point is most likely to have sprouted a rib whose endpoint is the shape point in question. (To stabilize the computation, we allow shape points to have mixed sources, treating them as probability-weighted mixtures of multiple ribs.) Part boundaries along the contour can be regarded as points at which responsibility for contour points switches from one axis to another (e.g., the boundaries between color-coded regions in the hand in Fig. 2; see below). The shape likelihood depends on this hypothesized ensemble of responsibilities, whereas the responsibilities depend on the currently estimated skeleton, suggesting a process similar to the well known expectation-maximization procedure, in which we alternately (i) estimate the correspondences (i.e., the ribs) between axial and contour points (the expectation phase), and (ii) search through the parametric space of skeletons, attempting to increase the posterior (decrease the DL) given the currently hypothesized correspondence (the maximization phase). This procedure is described in more detail in *Methods*.

Results

Figs. 3–5 show typical examples of the MAP skeleton, along with a conventional Voronoi-based implementation of Blum’s MAT (1, 15) shown in Figs. 3–5 *Insets* for comparison. Simple shapes (Fig. 3*a*) yield intuitive results devoid of spurious branches, and the estimated skeleton is robust against contour noise (Fig. 3*b* and *c*). Fig. 4 more specifically illustrates the robustness of the MAP skeleton as contour noise is introduced; the axial structure of the human form is recovered in a substantially invariant way in all three versions (*a*: no noise; *b*: noise throughout; *c*: noise on one arm and one leg only). Fig. 4*c* exemplifies the difficult case in which noise is added to some parts but not others, as in Richards *et al.*’s (28) famous “fuzzy pear,” which cannot be correctly handled by uniform smoothing techniques. Finally, Fig. 5 shows results for a variety of animal shapes. In each case the MAP skeleton corresponds closely to the intuitive part structure of the shape. The perceptual naturalness of these computed skeletons can be taken as “instant psychophysics,” supporting our claim that the MAP skeleton corresponds reasonably well to psychological shape representations.

A critical component of MAP skeleton estimation is the evaluation of candidate axes for inclusion in the hypothesized skeleton. As noted above, traditional approaches to computing the MAT have suffered from the problem of spurious axial branches, interfering with what otherwise might be a desirable isomorphism between the branches of a skeleton and the natural parts of a shape. The Bayesian approach provides a tool for




Fig. 3. Estimated MAP skeletons for the three simple shapes, showing the absence of forking (*a*) and the invariance to contour noise (*b* and *c*). (*a*) Rectangle. (*b*) Notched rectangle. (*c*) Noisy articulated blob. (*Insets*) Conventional Voronoi-based MAT.

handling this problem: a principled estimate of the statistical “significance” or evidence in favor of an axial branch. The relevant comparison is between a skeletal hypothesis SKEL that does not include the axial branch C and an augmented hypothesis $\text{SKEL}' = \text{SKEL} + C$ that does include it (Fig. 6). Following Bayes, we adopt the axial branch C if the posterior with it is better than the posterior without it, i.e.

$$\frac{p([\text{SKEL}+C]|\text{SHAPE})}{p(\text{SKEL}|\text{SHAPE})} > 1. \quad [6]$$

This condition can be easily restated in terms of DL,

$$DL(\text{SKEL}|\text{SHAPE}) - DL([\text{SKEL}+C]|\text{SHAPE}) > 0, \quad [7]$$

meaning that we should adopt axis C if doing so results in a net reduction in complexity (DL).

The difference in DLs is sometimes referred to as the weight of evidence, in this case quantifying the degree to which the added descriptive accuracy (or goodness of fit) of the augmented skeletal description offsets the added complexity of the additional axis. The criterion is thus a principled one in that it accurately reflects whether the new part yields a net benefit given the assumptions underlying the prior and likelihood functions. Because the weight of evidence quantifies the strength of posterior belief in the candidate axis C , it may serve to quantify the perceptual “salience” of the corresponding part of the shape (29).




Fig. 4. Estimated MAP skeletons for three variants of a human form, showing the invariance of the MAP skeleton to noise. (a) Human form. (b) Human form with added contour noise (independent and identically distributed circular Gaussian deviates added to each vertex). (c) Human form with contour noise on one arm and one leg only [compare with Richards *et al.*'s (28) fuzzy pearl]. (*Insets*) Conventional Voronoi-based MAT.

More broadly, this argument raises the possibility that Bayesian skeleton estimation might subsume such well known determinants of parts as the minima rule (9) (that extrema of negative

contour curvature constitute likely part boundaries) and the shortcut rule (10) (which holds that part cuts tend to be of minimal length). In the Bayesian framework each of these principles emerges as a by-product of the posterior maximization, in that, although they play no overt role in the computation, they tend to be obeyed by MAP skeletal estimates. For example, MAP skeleton estimation does not in any way involve the computation of curvature nor the identification of curvature extrema. Yet the boundaries between the points stochastically generated by one axial branch and another automatically tend to lie in regions of negative contour curvature and in particular near points of negative minima; so MAP skeletal estimates tend to place part boundaries at those locations. The boundaries between color-coded rib regions in Fig. 2, which lie between the fingers of the hand, are examples. The thrust of this argument admittedly parallels traditional, nonstochastic arguments, in this case amounting to a stochastic generalization of the transversality principle (30). (Distinct axes tend under our prior to be transverse, so stochastically generated ribs from distinct axes tend to point in transverse directions, meaning that collections of ribs stemming from a distinct axes tend to abut each other at points of rapidly changing tangent direction, as in Figs. 2 and 6.) But the Bayesian approach unifies what is otherwise a heterogeneous collection of disparate part-determination rules (31, 32) under the umbrella of one overarching inferential goal, the maximization of the skeletal posterior. Moreover, it provides a unified account of both part boundaries (points along the contour at which the shape divides) and part cuts (the resulting divisions of the shape itself), which in standard accounts each require its own distinct systems of rules, but which in ours both fall out of the same computation. Most importantly, our account has the potential to explain the well known pattern of exceptions to standard rules, for example, part boundaries that fall in positive-curvature regions of the bounding contour (31) and negative minima that are not perceived as part boundaries (8, 10, 33). Each of the traditional rules holds under some circumstances and fails under others, but always in accord with the maximization of the posterior.

Discussion

Among the most promising recent approaches to computing the MAT are those based on stochastic methods (34–37), which decrease sensitivity to noise by including a random element in the computation. Like ours, these methods presume a partly random rather than perfectly symmetric correspondence between opposing contour points and/or minimize some energy functional that embodies an asymmetry cost. Nevertheless, like Blum’s (1), most of these methods are based on the notion of local symmetry, although treated stochastically rather than deterministically. In particular, the medial axis itself is defined geometrically as the locus of points equidistant between contour points that have been designated as corresponding. In this sense, these existing stochastic methods fall short of a fully Bayesian approach in several respects, most specifically in that they lack an overt probabilistic shape-generating model for the skeleton, and for the shape given a skeleton, and thus lack well defined priors and likelihood functions.

In contrast, our approach adopts a full-fledged inverse-probability conception, setting as the computational goal the identification of a skeletal model most likely to have generated the shape (38). This approach entails overtly decomposing the observed shape into a “signal” process (the skeleton) and a “noise” process (the stochastic growth process). Of course, visual features treated as noise in some contexts might best be treated as signal in others (e.g., the notch in Fig. 3*b*, which might in some contexts be better treated as a bite or “negative part”), suggesting that top-down factors may play a role as well. A Bayesian approach has the benefit that such knowledge, when




Fig. 5. Estimated MAP skeletons for a variety of animal shapes. (Insets) Conventional Voronoi-based MAT.

available, may be readily incorporated into the model. As our results suggest, however, a generic default generative model gives reasonably good results with bottom-up geometry alone.

The MAP skeleton should not be regarded as an attempt to compute the MAT *per se*, but rather to estimate a related but distinct skeletal structure, the generative skeleton. The MAP




Fig. 6. Schematic of the Bayesian posterior test for the statistical contribution of an axis. The axis C is accepted if the posterior $p(\text{SKEL} + C | \text{SHAPE})$ with the axis is higher than the posterior $p(\text{SKEL} | \text{SHAPE})$ without it. (a) Without the candidate axis, the variance in rib lengths is very high, because the single axis must account for all points on the contour. (b) With the axis added to the hypothesized skeleton, the variance within each axis's (i.e., like-colored) collection of ribs is smaller. But this advantage comes as the cost of increased complexity in the skeleton, which entails a lower prior. The posterior test evaluates whether the added descriptive accuracy (higher likelihood) offsets the increased complexity (lower prior). If it does, the axis satisfies the criterion and is "accepted." The difference in the two log posteriors (DLs), the weight of evidence for the axial part, provides a measure of its subjective salience.

skeleton is not, in principle, necessarily medial, though it tends to maximize mediality, but only in conjunction with other properties not present in the MAT, such as skeletal simplicity, and low variance in the rib lengths. The previously noted problems connecting the MAT to psychological perceptions of shape are widely regarded as intrinsic to its fixed geometric definition. By contrast the MAP skeleton represents a more abstract perceptual shape description, which (like the MAT) brings out axial structure, but (unlike the MAT) does so in such a way that is both perceptually plausible and, in the sense that we have posed the problem, inferentially optimal.

The main benefit of our approach is the intuitive skeletons that MAP skeleton estimation tends to produce, with each axis corresponding to one perceived “part,” even with substantial contour noise (Figs. 3–5). This approach allows a compact, low-dimensional, but intuitive representation of shape, with enormous potential applications for shape recognition (12, 13), computer-based indexing of shape databases (18), and understanding of the function of long-range connections in visual cortex (4, 39, 40).

Moreover, our approach offers a number of important technical tools not provided by other methods, including a principled measure of the statistical evidence in favor of an axis (the difference in the log posteriors with and without the axis), the maximum value of the posterior over the space of skeletons, which gives a measure of how well any skeletal description explains the given shape, and a principled measure of shape complexity, the DL of the MAP skeleton. Each of these quantities has a natural psychological correlate, respectively, the subjective part salience, the subjective “axiality” of the shape, and the subjective complexity of the shape, none of which has received rigorous definitions in the literature before to our knowledge. All of these advantages stem from the underlying idea of formulating shape representation as a Bayesian inference problem, bringing it into line with a growing segment of modern perceptual theory (20, 21) and drawing closer to Attneave’s original goal (41) of understanding shape as an information-processing problem.

Methods

Here, we sketch a computational procedure for estimating the MAP skeleton. As mentioned, we regard our theory as a “theory of the computation,” not a processing model; the computational implementation should be taken simply as a “proof of concept” that the MAP skeleton is computable and has the intended desirable properties and not as a realistic model of neural shape processing.

We seek the skeletal description with minimum DL, defined as the negative logarithm of the posterior probability $p(\text{SKEL}|\text{SHAPE})$. The rib length function is estimated by pooling ribs within a moving mask with a fixed width (plus or minus $\approx 1/3$ the length of a typical axis in the examples shown), enforcing the constraint of a continuous length function connected to each axis. (This pooling introduces some dependence among the estimated ribs, making Eq. 3 only an approximation.) We also assume a von Mises distribution on the deviation between the inward-pointing shape normal and the rib, which is amplified when this deviation exceeds $\pi/2$ and has the effect of preventing “explanation from outside the shape.”

To estimate the skeleton, we use a gradient descent procedure loosely based on expectation–maximization. We use the conventional Voronoi-based MAT (15) to form an initial, grossly overfitted estimate of the skeleton. This point set is organized into a hierarchical structure by merging axes so as to maximize collinearity within each axis. Then all nonroot axes are subjected to the Bayesian posterior ratio test of significance (Eq. 4); axes failing the test are then pruned. The remaining axes are parameterized by using a piecewise cubic spline approximation, with knot points at every axial branch point, and additional knot points chosen successively until the spline approximation fits the original axis to within a fixed tolerance, resulting in a variable number m of knot points per axis. This procedure yields a representation having $2m$ parameters per axis (plus one additional parameter required to code the location of the root axis). With this parametric description as a starting point, an iterative gradient procedure is initiated, with two stages alternating:

1. Estimate the “ribs” by associating each contour point with some set of axis points that explain it. For each shape point x , we choose the axis point and side (left or right) that assigns x the highest likelihood (Eq. 3).
2. With the rib correspondences fixed, take one step down the gradient of DL (equivalently, up the gradient of posterior). We use a standard simplex method (Nelder–Mead) to execute the gradient descent.

This procedure converges to an estimate of the MAP skeleton, examples of which are shown in Figs. 4 and 5.

We thank Elan Barenholtz, Randy Gallistel, and Eileen Kowler for helpful comments. J.F. was supported by National Science Foundation Grant 9875175 and National Institutes of Health Grant R01 EY15888, and M.S. was supported by National Science Foundation Grant BCS-021694. Shapes in Figs. 2b, 4a, and 5 were obtained from the Laboratory for Engineering Man/Machine Systems at Brown University (Providence, RI).

1. Blum H (1973) *J Theor Biol* 38:205–287.
2. Blum H, Nagel RN (1978) *Pattern Recogn* 10:167–180.
3. Lee TS (2003) *J Physiol (Paris)* 97:121–139.
4. Lee TS, Mumford D, Romero R, Lamme VAF (1998) *Vision Res* 38:2429–2454.
5. Kovacs I, Julesz B (1994) *Nature* 370:644–646.
6. Kovacs I, Feher A, Julesz B (1998) *Vision Res* 38:2323–2333.
7. Burbeck CA, Pizer S (1995) *Vision Res* 35:1917–1930.
8. Barenholtz E, Feldman J (2003) *Vision Res* 43:1655–1666.
9. Hoffman DD, Richards WA (1984) *Cognition* 18:65–96.
10. Singh M, Seyranian GD, Hoffman DD (1999) *Percept Psychophys* 61:636–660.
11. Marr D, Nishihara HK (1978) *Proc R Soc London Ser B* 200:269–294.
12. Biederman I (1987) *Psychol Rev* 94:115–147.
13. Tarr MJ, Bulthoff HH, Zabiniski M, Blanz V (1997) *Psychol Sci* 8:282–289.
14. Kimia BB, Tannenbaum AR, Zucker SW (1995) *Int J Comput Vision* 15:189–224.
15. Ogniewicz RL, Kubler O (1995) *Pattern Recogn* 28:343–359.
16. Geiger D, Liu T-L, Kohn RV (2003) *IEEE Trans Pattern Analysis Machine Intelligence* 25:86–99.
17. Katz RA, Pizer SM (2003) *Int J Comput Vision* 55:139–153.
18. Siddiqi K, Shokoufandeh A, Dickinson S, Zucker S (1999) *Int J Comput Vision* 30:1–24.
19. Leyton M (1989) *Cognit Sci* 13:357–387.
20. Knill D, Richards W, eds (1996) *Perception as Bayesian Inference* (Cambridge Univ Press, Cambridge).
21. Kersten D, Mamassian P, Yuille A (2004) *Annu Rev Psychol* 55:271–304.
22. Geisler WW, Kersten D (2002) *Nat Neurosci* 5:508–510.
23. Feldman J, Singh M (2005) *Psychol Rev* 112:243–252.
24. Mardia KV (1972) *Statistics of Directional Data* (Academic, London).
25. Feldman J (2001) *Percept Psychophys* 63:1171–1182.
26. Thompson DW (1942) *On Growth and Form* (Cambridge Univ Press, Cambridge).
27. Rissanen J (1989) *Stochastic Complexity in Statistical Inquiry* (World Scientific, Singapore).
28. Richards W, Dawson B, Whittington D (1986) *J Opt Soc Am A* 3:1483–1491.
29. Hoffman DD, Singh M (1997) *Cognition* 63:29–78.
30. Bennett B, Hoffman D (1987) in *Image Understanding*, ed Richards WA (Ablex, Norwood, NJ), pp 215–256.
31. Singh M, Hoffman DD (2001) in *From Fragments to Objects: Segmentation and Grouping in Vision: Advances in Psychology*, eds Shipley T, Kellman P (Elsevier, New York), Vol 130, pp 401–459.
32. de Winter J, Wagemans J (2006) *Cognition* 99:275–325.
33. Siddiqi K, Tresness KJ, Kimia BB (1996) *Perception* 25:399–424.
34. Tsao Y-F, Fu K-S (1984) *Comput Vision Graphics Image Processing* 25:348–370.
35. Zhu SC, Yuille AL (1996) *Int J Comput Vision* 20:187–212.
36. Zhu S-C (1999) *IEEE Trans Pattern Anal Mach Intell* 21:1158–1169.
37. Kegl B, Krzyzak A (2002) *IEEE Trans Pattern Anal Mach Intell* 24:59–74.
38. Telea A, Sminchisescu C, Dickinson S (2004) in *Proceedings of the International Conference on Pattern Recognition* (IEEE Computer Society, Los Alamitos, CA), Vol 4, pp 19–22.
39. Gilbert CD (1995) in *The Cognitive Neurosciences*, ed Gazzaniga MS (MIT Press, Cambridge, MA), pp 73–90.
40. Spillman L, Werner JS (1996) *Trends Neurosci* 19:428–434.
41. Attneave F (1954) *Psychol Rev* 61:183–193.