

Bayesian contour integration

JACOB FELDMAN

Rutgers University, New Brunswick, New Jersey

The process by which the human visual system parses an image into contours, surfaces, and objects—*perceptual grouping*—has proven difficult to capture in a rigorous and general theory. A natural candidate for such a theory is Bayesian probability theory, which provides optimal interpretations of data under conditions of uncertainty. But the fit of Bayesian theory to human grouping judgments has never been tested, in part because methods for expressing grouping hypotheses probabilistically have not been available. This paper presents such methods for the case of *contour integration*—that is, the aggregation of a sequence of visual items into a “virtual curve.” Two experiments are reported in which human subjects were asked to group ambiguous configurations of dots (in Experiment 1, a sequence of five dots could be judged to contain a “corner” or not; in Experiment 2, an arrangement of six dots could be judged to fall into two disjoint contours or one smooth contour). The Bayesian theory accounts extremely well for subjects’ judgments, explaining more than 75% of the variance in both tasks. The theory thus provides a far more quantitatively precise account of human contour integration than has been previously possible, allowing a very precise calculation of the subjective goodness of a virtual chain of dots. Because Bayesian theory is inferentially optimal, this finding suggests a “rational justification,” and hence possibly an evolutionary rationale, for some of the rules of perceptual grouping.

Perceptual grouping is the process whereby individual items in the visual image are aggregated into larger structures. Grouping is known to influence many low-level visual computations, such as the perception of lightness (Adelson, 1993; Gilchrist, 1977), the perception of motion (Shimojo & Nakayama, 1990; Weiss & Adelson, 1998), and visual search (He & Nakayama, 1992; Treisman, 1982). Yet the process by which a grouping interpretation is chosen, often described in terms of somewhat vague and poorly understood Gestalt principles, has proven difficult to characterize precisely. Perhaps the main obstacle has been the difficulty in specifying in a mathematically rigorous way the various candidate interpretations from which the visual system unconsciously chooses, and the function determining subjective preference among these interpretations. This paper attempts to develop such a theory in the specific case of *contour integration*—that is, the aggregation of a sequence of visual items into a *virtual curve*—and then to test the predictions of this model against the judgments of human observers.

Bayesian Approaches to Perception

A natural candidate for a rigorous model is Bayesian probability theory, which has often been advocated as an optimal method for making decisions under conditions of

uncertainty (Jaynes, 1983) and has recently attracted a great deal of interest among investigators of human vision (Bülthoff & Yuille, 1991; Knill & Richards, 1996; Landy Maloney, Johnston, & Young, 1995). Bayesian or quasi-Bayesian models have been brought to bear on the interpretation of motion (Weiss & Adelson, 1998) and surfaces (Mamassian & Landy, 1998; Nakayama & Shimojo, 1992), recognition of objects (Liu, Knill, & Kersten, 1995), classification of shapes (Feldman, 2000), and combination of distinct cue sources (Landy et al., 1995; Yuille & Bülthoff, 1996). The application of Bayesian theory to grouping may be more difficult than in these other cases, because, in grouping, the target inference—the “best grouping”—is difficult to describe formally and, arguably, might not admit an objective definition (this issue is discussed below). A fully realized Bayesian theory of human perceptual grouping would need to spell out the observer’s subjective model of what alternative grouping hypotheses are possible, how they might give rise to possible image configurations, and with what likelihoods.

In Bayesian theory, the degree of belief in a perceptual hypothesis H_0 (henceforth called the *target hypothesis* or *target interpretation*) given image I is expressed by the *posterior probability*:

$$p(H_0 | I) = \frac{p(I | H_0)p(H_0)}{\sum_i p(I | H_i)p(H_i)}, \quad (1)$$

where H_0, H_1, \dots are candidate interpretations, $p(H_i)$ is the prior probability of hypothesis H_i , and $p(I | H_i)$, called a *likelihood* term, is the probability that the observed image I would be generated by the hypothesis H_i . The likelihood term is a measure of fit between the hypothesis under consideration and the image configuration. The prospect

This work was supported by National Science Foundation Grant SBR-9875175. I am grateful to G. John Andersen and two anonymous reviewers for helpful comments on the manuscript and to Henry Chi for assistance in data collection. Address correspondence to J. Feldman, Department of Psychology, Center for Cognitive Science, Rutgers University, Busch Campus, New Brunswick, NJ 08903 (e-mail: jacob@ruccs.rutgers.edu).

of formulating a Bayesian model of grouping hinges on the construction of suitable likelihood terms for grouping hypotheses.

We focus on the case of grouping individual visual items (dots, edge fragments, etc.) into smooth contours, a process known to occur early and to be essential in the construction of visual representations (Caelli & Uman-sky, 1976; Glass, 1969). The visual system’s tendency to extract approximately collinear patterns from the image has been investigated in some detail (Feldman, 1996, 1997; Pizlo, Salach-Golyska, & Rosenfeld, 1997; Smits & Vos, 1987; Smits, Vos, & van Oeffelen, 1985). Yet, there is still no quantitative model that will predict both (1) the subjective coherence of a dot pattern as a function of its geom-etry and (2) the particular grouping interpretation that a human observer will perceive in an ambiguous configura-tion (e.g., the particular assignment of dots to distinct virtual curves. Moreover, what is known about the quantita-tive properties of curve grouping does not afford any convenient mathematical generalization to other types of grouping problems, such as grouping into surfaces and ob-jects. Such a generalization might be provided by Bayesian theory, which is in principle completely general in its ap-plication.

A major obstacle is the lack of a model for how the sys-tem combines the many local estimates of collinearity (e.g., the outputs of local orientation-tuned cells) into a single global judgment of curve coherence, a problem some-times referred to as *cooperativity* (Kubovy & Wagemans, 1995; Zucker, Stevens, & Sander, 1983). It is believed that raw judgments of collinearity propagate laterally in visual cortex (Field, Hayes, & Hess, 1993), but the mathematical form of the combination rule is unknown. Arguments from differential geometry suggest that along a smooth curve, sampled at intervals to produce visible points, successive angles between points should tend to be collinear, and the implicit curve should be well approximated by the local tangent (Parent & Zucker, 1989). But such arguments do not specify exactly how much deviation from collinearity should suppress the impression of a subjective curve, nor how successive angles should interact (i.e., the combina-tion rule). These lacunae need to be repaired in order to construct suitable subjective likelihood functions.

A Bayesian Model of Smooth Curves

Earlier studies (Feldman, 1996, 1997) have suggested a mathematical form for the likelihood function corre-sponding to subjectively smooth curvilinear patterns. The simplest case is three visual items, parameterized by an angle α_1 measuring the deviation from perfect collinear-ity (0°) (see Figure 1). Human judgments of apparent curvilinearity are consistent with a model in which,

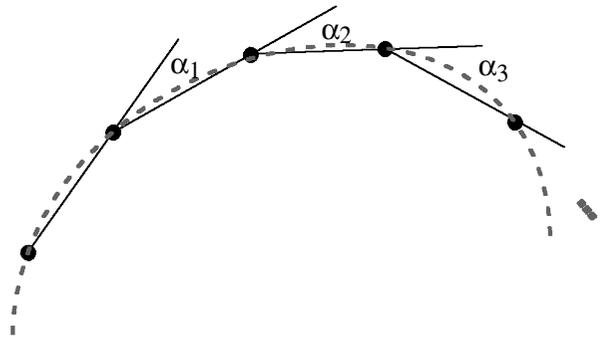


Figure 1. Notation of angles along a curvilinear pattern of dots.

under the hypothesis of a smooth curve, the expected distribution of angle α_1 is proportional to a Gaussian dis-tribution centered at 0° . That is, the likelihood function for the smooth hypothesis, $p(\alpha_1|\text{smooth})$, is given by

$$L_3(\alpha_1) = h_3 \exp\left[-\frac{1}{2}(\alpha / \sigma_3)^2\right], \tag{2}$$

where σ_3 is the standard deviation of the Gaussian dis-tri-bution, and h_3 is a proportionality constant. In the case of four visual items, there are now two successive angles, α_1 and α_2 , and the joint distribution $p(\alpha_1, \alpha_2|\text{smooth})$ is given by

$$L_4(\alpha_1, \alpha_2) = h_4 \exp\left\{-\frac{1}{2(1-r^2)}\left[\left(\frac{\alpha_1}{\sigma_4}\right)^2 + \left(\frac{\alpha_2}{\sigma_4}\right)^2 - 2r\left(\frac{\alpha_1\alpha_2}{\sigma_4^2}\right)\right]\right\}, \tag{3}$$

where σ_4 is the standard deviation of each of the two mar-ginal distributions, r is the correlation coefficient, and h_4 is a proportionality constant (Figure 2). Empirical esti-mates of correlation r have shown it to be nonzero between successive angles—that is, successive angles along a sub-jectively smooth curve are not independently distributed—but, approximately zero between nonsuccessive angles (Feldman, 1997). This means that by Bayes’ rule, the like-lihood function for the general case of n items is the prod-uct of successive iterations of the function L_4 (Equation 4).

Equation 4 represents a “moving four-item window” operating on successive angle pairs, each of which con-tributes independently to the overall perception of a smooth curve (Figure 3). In the spectrum of Bayesian mod-els, the assumption of correlation between successive an-gles but independence between nonsuccessive angles

$$p(\alpha_1, \alpha_2, \dots, \alpha_n | \text{smooth}) = L_4(\alpha_1, \alpha_2)L_4(\alpha_2, \alpha_3)\dots L_4(\alpha_{n-1}, \alpha_n) \tag{4}$$

$$= \prod_{i=1}^{n-1} L_4(\alpha_i, \alpha_{i+1})$$

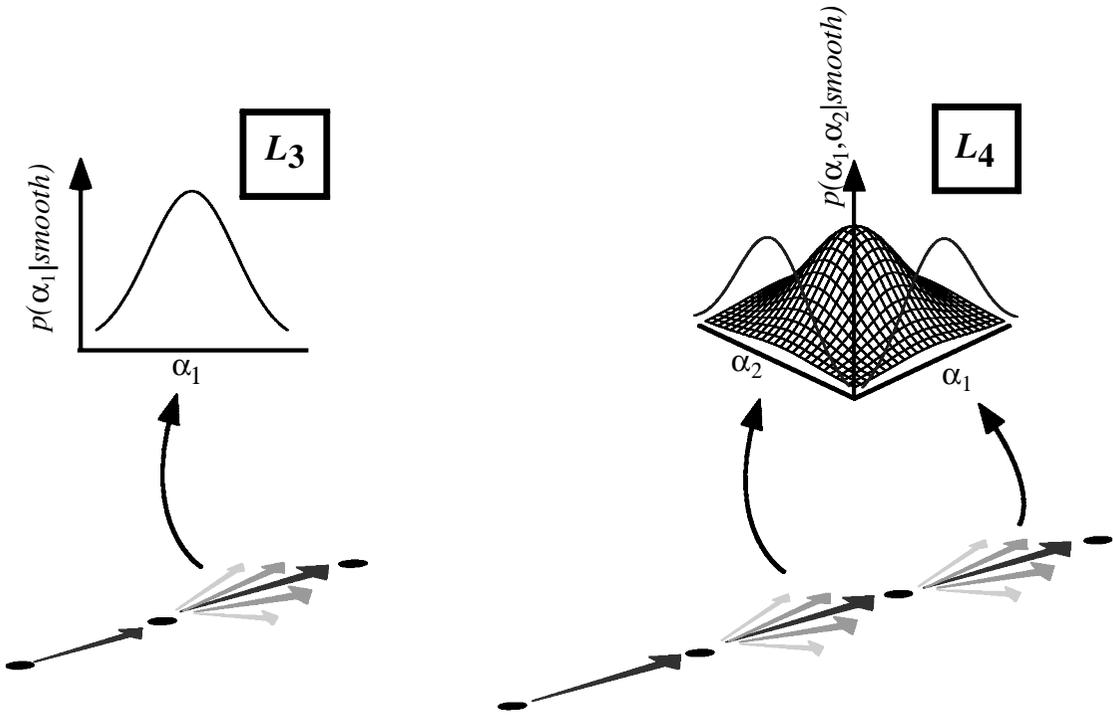


Figure 2. The atomic functions L_3 and L_4 , used in the construction of Bayesian models. The functions give the expected distribution of angles along a subjectively smooth virtual curve.

places this model somewhere between *weak fusion* (all cues are assumed to be independent) and *strong fusion* (all joint densities are computed) (Yuille & Bülhoff, 1996), approximately in the style of *modified weak fusion* (Landy et al., 1995).

The two “atomic” likelihood functions L_3 and L_4 together can be used to construct probabilistic models of arbitrarily long smooth curves (L_3 is needed only when a curve contains only three items; the exact procedure is detailed below), and, moreover, to build complete scene representations consisting only of piece-wise smooth contours. Hence, the resulting composite probabilistic functions provide rigorous numerical models of how well candidate grouping interpretations fit the observed configuration, allowing the visual system to in effect select the probabilistically optimal grouping interpretation. It is worth noting that, notwithstanding their superficially complex mathematical form, these functions may easily be computed by simple arrangements of neural hardware (Feldman, 1997).

The experiments reported below investigate human subjects’ subjective grouping of dot configurations into piece-wise smooth virtual curves. It is to be emphasized that the Bayesian models presented below as accounts of subjects’ data are constructed entirely out of the atomic functions L_3 and L_4 and contain no unmotivated or ad hoc components.

Experiments

Two types of tasks were employed. In the corners task, displays contained five dots parameterized by three an-

gles, a_1 , a_2 , and a_3 (Figure 4, top). In the corners task, the angles used were 0° , $\pm 15^\circ$, $\pm 30^\circ$, and $\pm 45^\circ$ (a_1 , a_3), and 0° , $\pm 10^\circ$, $\pm 20^\circ$, $\pm 30^\circ$, $\pm 40^\circ$, $\pm 50^\circ$, and $\pm 60^\circ$ (a_2), all fully crossed, for a total of $7 \times 13 \times 7 = 637$ combinations. As illustrated in Figure 4, these parameters allow for a wide variety of configurations, ranging from some that clearly appear to have a corner, to some that appear to be quite smooth. In the two-contours task, angles were 0° , $\pm 15^\circ$, $\pm 30^\circ$, and $\pm 45^\circ$ (a_1 , a_3), and 0° , $\pm 10^\circ$, $\pm 20^\circ$, $\pm 40^\circ$, and $\pm 60^\circ$ (a_2), all fully crossed, for a total of $7 \times 7 \times 7 = 343$ combinations. Again, these parameters allow for a wide range of configurations (Figure 4), including some that appear to spontaneously “break” into two contours, eliciting a *two contour* response, as well as some that suggest a single smooth contour.

Dots were dark circular patches (subtending 0.11° of visual angle in the corners task, and 0.055° or 0.11° in the two-

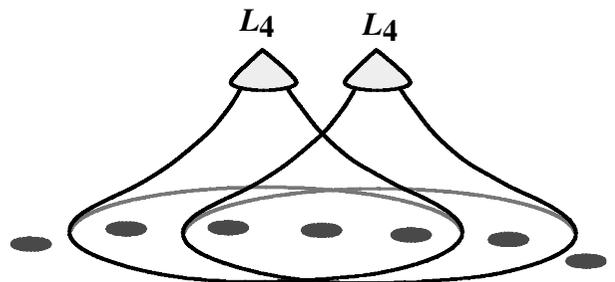


Figure 3. L_4 is computed in parallel on groups of four dots lying successively along a chain of dots.

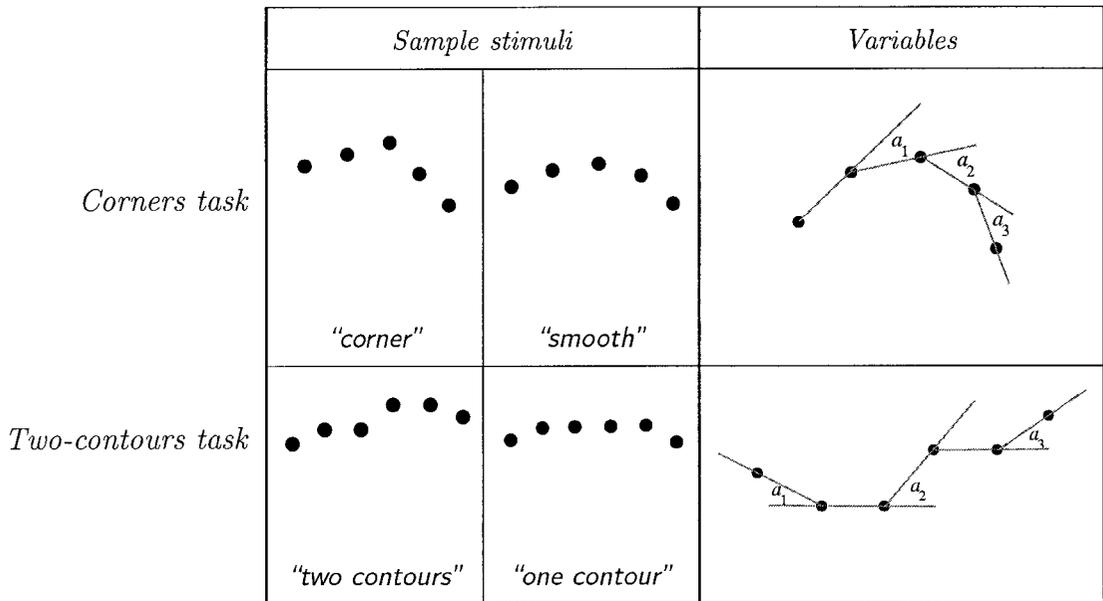


Figure 4. The two tasks employed, showing sample stimuli (left, with likely responses; stimuli not drawn to scale) and illustration of the experimental variables a_1 , a_2 , and a_3 (right).

contours task) on a uniform white background displayed at high contrast in a darkened room at a 60-cm viewing distance, with observers' heads fixed by a chinrest.

In the corners task, each configuration was displayed in a randomly chosen orientation. In the two-contours task, configurations were presented upright as in Figure 4. (That is, each figure was displayed so that the second and third dots were at the same height as each other, and likewise the fourth and fifth dots.) Nineteen naive subjects were asked to judge, on a 1–5 scale, whether the dots traced out a corner or a single smooth curve (1 = *definitely a smooth curve*, 5 = *definitely a corner*). In the two-contours task, displays contained six dots again parameterized by three angles, a_1 , a_2 , and a_3 (Figure 4, bottom). (In addition, two stimulus sizes were employed, but no scale effects were found, and henceforth the data are presented collapsed across scale.) Seventeen naive subjects, none of whom had participated in the corners experiment, were asked to rate whether the display contained two distinct smooth contours or one long smooth contour, again on a 1–5 scale (1 = *definitely one smooth contour*, 5 = *definitely two smooth contours*). Subjects' mean ratings of each condition, after normalization to the interval (0,1), were taken to represent the subjects' a posteriori belief that the stimulus configuration belonged to the target interpretation.

The two tasks were chosen in order to reflect two fundamental modes of contour extraction: Dots can be assigned to two completely disjoint contours, or they can be assigned to two distinct sections of the same contour that are separated by a perceived tangent discontinuity (Link & Zucker, 1987). One of the advantages of a Bayesian approach is the possibility of treating these two modes of

grouping in a theoretically uniform manner, and, in fact, both Bayesian models described below (one for each task) draw on the same probabilistic vocabulary—namely, the functions L_3 and L_4 .

Bayesian Models of the Two Tasks

In each task, several grouping interpretations are possible, some leading to the target interpretation (*corner* or *two contours*, respectively), others leading to the perception of a single smooth contour. One immediate complication is that in each task, there are several different perceptually distinct interpretations that all lead to the target response. Consider first the corners task. In this task, one may perceive a corner at the central dot; denote this interpretation by H_c . Alternatively, one may perceive a corner at either the second or the fourth dot, again leading to a *corner* response; denote these interpretations by $H_{c'}$ and $H_{c''}$. Counterposed to these is the *smooth* interpretation H_s . All hypotheses under consideration are depicted schematically in Figures 5 (corners) and 6 (two-contours).

For any hypothesis H_i , denote by P_i the product $p(H_i)p(I|H_i)$ of its prior and its likelihood. By Bayes' rule, the probability of the target response (*corner*, regardless of where the corner is perceived) is

$$p(\text{corner} | I) = h \frac{P_c + P_{c'} + P_{c''}}{P_c + P_{c'} + P_{c''} + P_s}, \quad (5)$$

where h is a free scaling factor relating this expression to the subjects' numeric ratings. Likelihood functions $p(I|H_i)$ were constructed for each interpretation in the following manner. Three dots at an angle, α_1 , are assigned likelihood $L_3(\alpha_1)$. Four or more dots with angles $\alpha_1, \alpha_2, \dots$ are assigned likelihood by concatenations of the function L_4 as

in Equation 4. One- or two-dot groups are each perfectly consistent with a straight line and hence have likelihood unity; they drop out of the resulting formulae. The full Bayesian model for the corners task is then provided by Equation 5, substituting $p(H_i)p(I|H_i)$ for each P_i , and then using Figure 5 to provide expressions for each likelihood term $p(I|H_i)$.

The free parameters of the Bayesian model include three parameters, σ_3 , σ_4 , and r , of the atomic functions L_3 and L_4 , the overall scaling parameter h , and the priors. Although there are four separate scalar priors in Equation 5 (one for each hypothesis), there are in fact really only two degrees of freedom among the priors, after one assumes $p(H_c) = p(H_{c'})$ (by symmetry) and further expresses all the priors relative to one standard prior chosen arbitrarily. [In the analysis, $p(H_s)$ is omitted, implicitly representing $p(H_c)$ and $p(H_{c'})$ as proportions of it.] Boiling this all down, the Bayesian model for the corners task contains six free parameters: σ_3 , σ_4 , r , $p(H_c)$, $p(H_{c'})$, and h . The first three are the parameters of the atomic likelihood functions L_3 and L_4 ; the next two are the free priors, and the last is the overall scaling factor. The free parameters are admittedly more numerous than in some previous contour integration theories, but all are motivated directly by Bayesian theory and readily admit meaningful interpretation. In the analysis below, these six variables are treated as free parameters in a nonlinear regression fitting the Bayesian model to subjects' ratings.

An extremely similar analysis applies to the two-contours tasks, with hypotheses H_2 , $H_{2'}$, and $H_{2''}$ associated

with the response *two contours*, and H_s associated with *one contour* (Figure 6; note that here H_s has a different mathematical form than in the corners task due to the different stimulus geometry). Figures 5 and 6 give explicit expressions for the likelihood of each hypothesis under consideration in both tasks.

Results

First, for both tasks, the effects and interactions of all three angular variables, a_1 , a_2 , and a_3 , were submitted to an analysis of variance (ANOVA). In both tasks, all three main effects, all 3 two-way interactions, and the three-way interaction were significant at $p < .0001$ (statistical details are given in Table 1). Figures 7 (corners) and 8 (two-contours) show the main effects of a_1 , a_2 , and a_3 (along with the Bayesian model, discussed below), and Figures 9 and 10 show the 3 two-way interactions $a_1 \times a_2$, $a_1 \times a_3$, and $a_2 \times a_3$. The most salient main effects were that, in both tasks, target interpretations (1) increased markedly as angle a_2 increased and (2) decreased as a_1 and a_2 increased, except at the tails, where target interpretations again increased. The effect of a_2 was much larger in magnitude than that of a_1 and a_3 .

The significant interactions suggest a nonlinear decision surface, and, indeed, the plots of the 3 two-way interactions (Figures 9 and 10) show highly curved surfaces. As remarked by Jaynes (1993), "Bayes' theorem automatically generates the exact nonlinear function called for by the problem" (p. 268), and hence it might be hoped that Bayesian theory would provide a quantitative account of

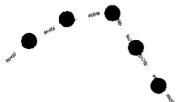
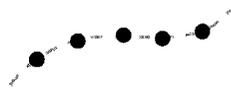
Hypothesis	Likelihood	
H_c	 $L_3(a_1)L_3(a_3)$	} "corner"
$H_{c'}$	 $L_4(a_1, a_2)$	
$H_{c''}$	 $L_4(a_2, a_3)$	
H_s	 $L_4(a_1, a_2)L_4(a_2, a_3)$	"smooth"

Figure 5. Candidate hypotheses in the corners task, showing illustration (left) and mathematical form (right).

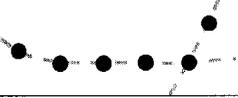
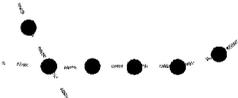
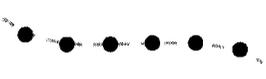
<i>Hypothesis</i>	<i>Likelihood</i>	
H_2 	$L_3(a_1)L_3(a_3)$	} "two contours"
$H_{2'}$ 	$L_4(a_1, a_2)L_4(a_2, -a_2)$	
$H_{2''}$ 	$L_4(a_2, -a_2)L_4(a_2, a_3)$	
H_s 	$L_4(a_1, a_2)L_4(a_2, -a_2)L_4(a_2, a_3)$	"one contour"

Figure 6. Candidate hypotheses in the two-contours task, showing illustration (left) and mathematical form (right).

the shapes of these surfaces. Hence, in the next analysis, the Bayesian model derived above was fit to the full four-dimensional decision surface (probability as a function of $a_1, a_2,$ and a_3).

Figures 7 and 8 show the best-fit Bayesian model (chosen by Levenburg–Marquardt, using least-squared error) superimposed on the subjects’ data. For ease of viewing, the Bayesian model is shown superimposed on the marginal means (main effects) only, but note that the model shown reflects a fit not just to this relatively small number of data points, but rather to the full 4-D response surface, comprising 637 independent data points in the corners task, and 343 in the two-contours task, while using only six degrees of freedom in each model. The fit is extremely good [corners, $R^2 = .8443$; $F(6,631) = 88.78, p < .000001$; two-contours, $R^2 = .7686$; $F(6,337) = 43.17307, p < .000001$], although in the two-contours data, the subjects’ responses seem slightly more peaked than in the

model. The good fit of the models to relatively subtle effects (e.g., the rising tails in a_1 and a_3) reflects the power of Bayesian machinery to make quantitative predictions beyond what is intuitively obvious. In particular, the Bayesian model makes an accurate prediction about exactly how various angular cues are combined—the problem of *cooperativity* that is rarely successfully addressed in conventional theories.

Estimated parameters of the Bayesian models are given in Table 2. All six parameters in each model made significant contributions to the fit, as determined by comparing fits with versus without each parameter (details are given in Table 2). The parameters $\sigma_3, \sigma_4,$ and r controlling the atomic functions L_3 and L_4 are very similar in the two tasks (see the first three rows in Table 2). This result is especially crucial because it supports the hypothesis that, notwithstanding the difference between the two tasks, subjects’ judgments were constructed from a fixed and stable probabilistic lexicon. It is worth noting, however, that these estimates differ from earlier estimates (Feldman, 1997). These differences may be attributable to the difference in tasks (judging the smoothness of a single dot pattern vs. grouping an entire configuration into separate smooth components), although this is admittedly speculative. Such context-driven modulation of parameters may be related to the neurophysiological finding of rapid modulation of lateral connections within visual cortex (Gilbert, 1995) and, in Bayesian theory, to the mathematical notion of *competitive priors* (Yuille & Bülthoff, 1996).

Table 1
Details of the Analyses of Variance

Effect	Corners zTask		Two-Contours Task	
	<i>F</i>	<i>p</i> <	<i>F</i>	<i>p</i> <
a_1	6.141	.0001	25.486	.0001
a_2	502.242	.0001	1095.154	.0001
a_3	5.778	.0001	10.056	.0001
$a_1 \times a_2$	9.955	.0001	27.971	.0001
$a_1 \times a_3$	2.882	.0001	26.278	.0001
$a_2 \times a_3$	9.140	.0001	22.679	.0001
$a_1 \times a_2 \times a_3$	1.701	.0001	1.620	.0001

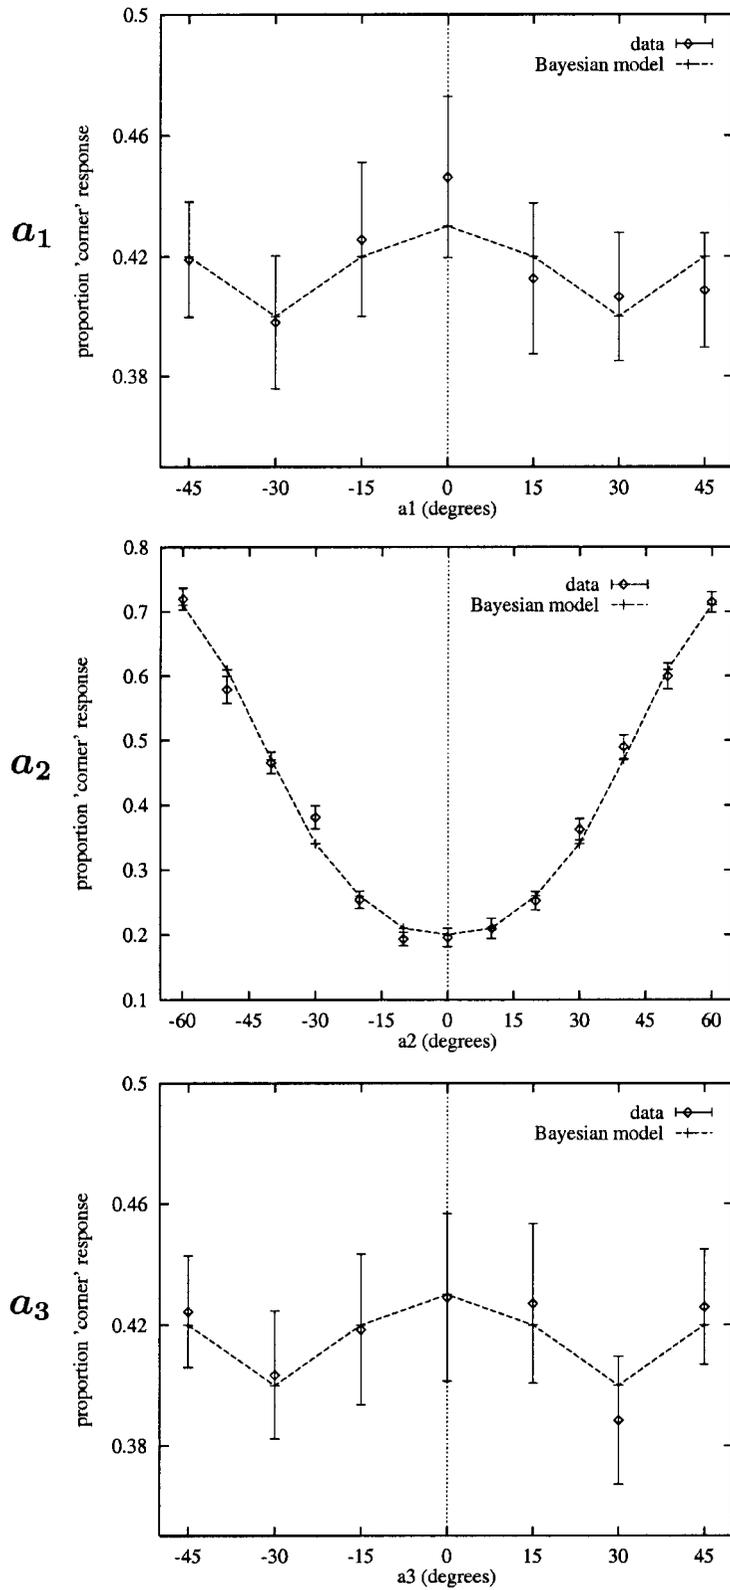


Figure 7. Subjects' data and the Bayesian model in the corners task, showing proportion *corners* judgment as a function of a_1 (top), a_2 (middle), and a_3 (bottom). Error bars are standard errors.

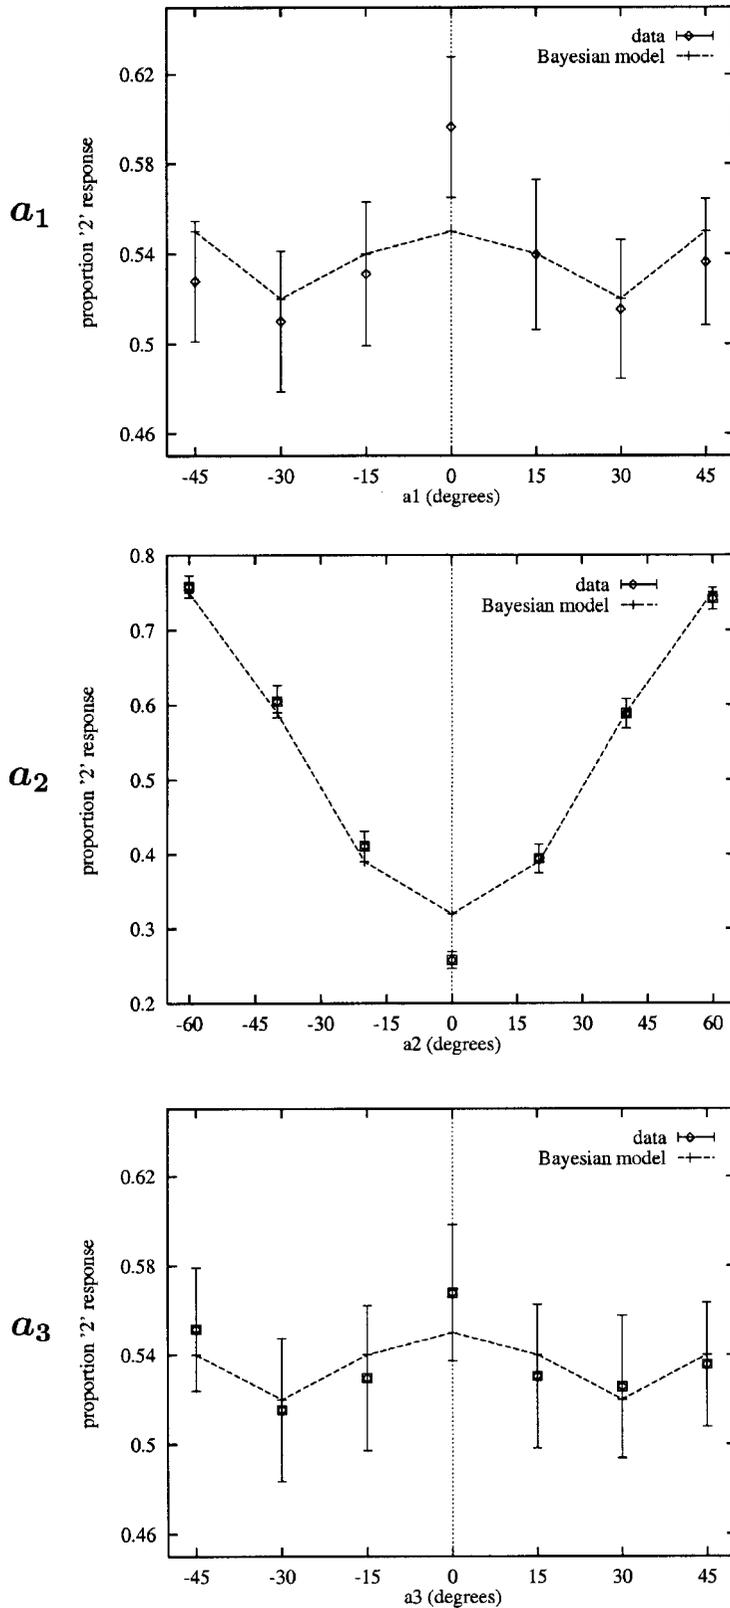


Figure 8. Subjects' data and the Bayesian model in the two-contours task, showing proportion *two contours* judgments as a function of a_1 (top), a_2 (middle), and a_3 (bottom). Data shown are collapsed over scale. Error bars are standard errors.

Conclusions

The very close fit of the Bayesian model compares favorably with the results from existing theories in other grouping domains. By comparison, only a few years ago, numerical theories of proximity grouping were judged to be empirically successful simply when they agreed with human judgments more often than chance (Compton & Logan, 1993), although more recently, standards of empirical success in that area—but not in collinearity grouping—have improved (Kubovy, Holcombe, & Wagemans, 1998). This good fit means that with the aid of Bayesian calculations, the subjective goodness of a virtual curve of dots can be calculated with about an order of magnitude more precise than with conventional treatments. (In most cur-

rent treatments, curves are classified simply as *curved* [angles generated at random from, e.g., $\pm 30^\circ$] or *straight* [angles near 0°]; this means that subjective goodness is treated as if it depended only on the mean angle over the entire curve, a far coarser measure than the combinations of L_4 used here.)

Some of the goodness of fit could be due to the choice of Gaussian likelihood functions and the use of pooled data, because sums of independent distributions tend toward normality (the central limit theorem). However, statistical analysis of individual subject's data in earlier studies (Feldman, 1997) showed that the model fit is reduced only slightly, and visual inspection of individual subjects' data in the present experiments suggests the same. In any

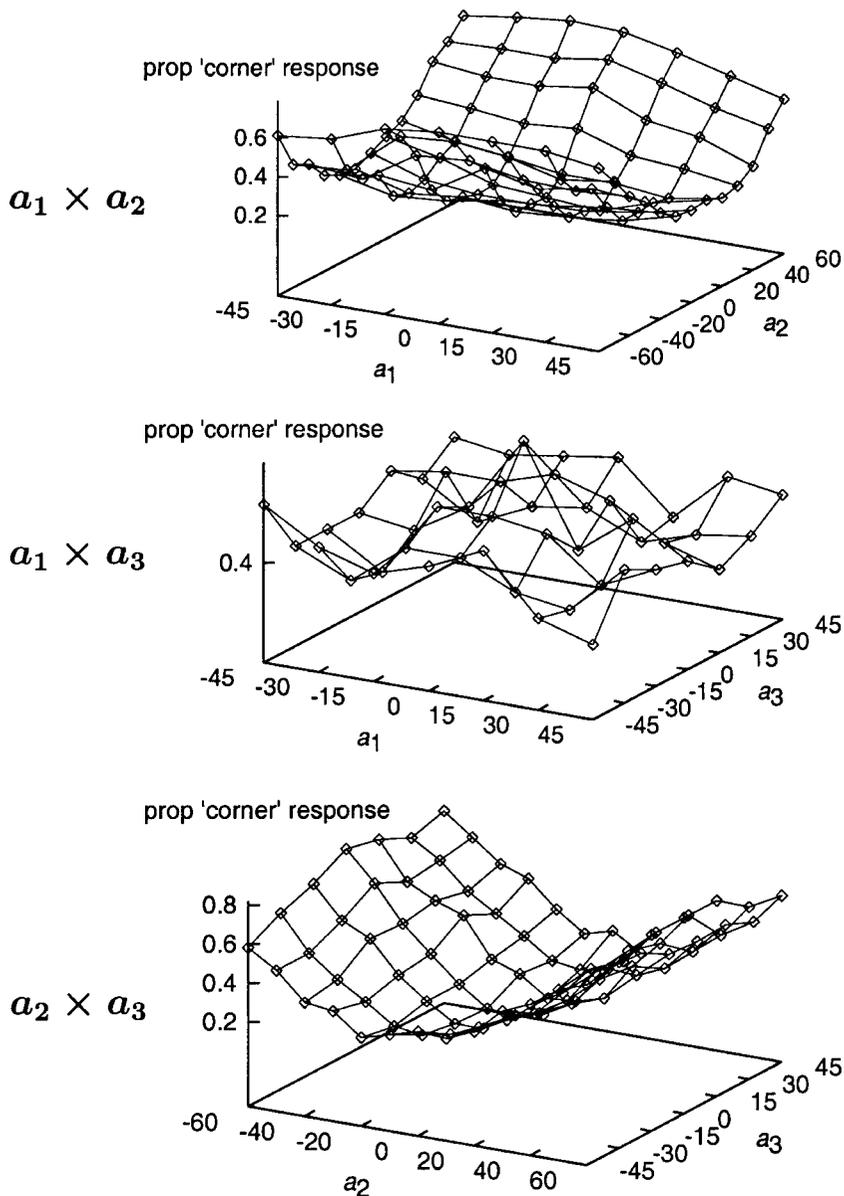


Figure 9. Two-way interactions in the corners task: $a_1 \times a_2$ (top), $a_1 \times a_3$ (middle), and $a_2 \times a_3$ (bottom).

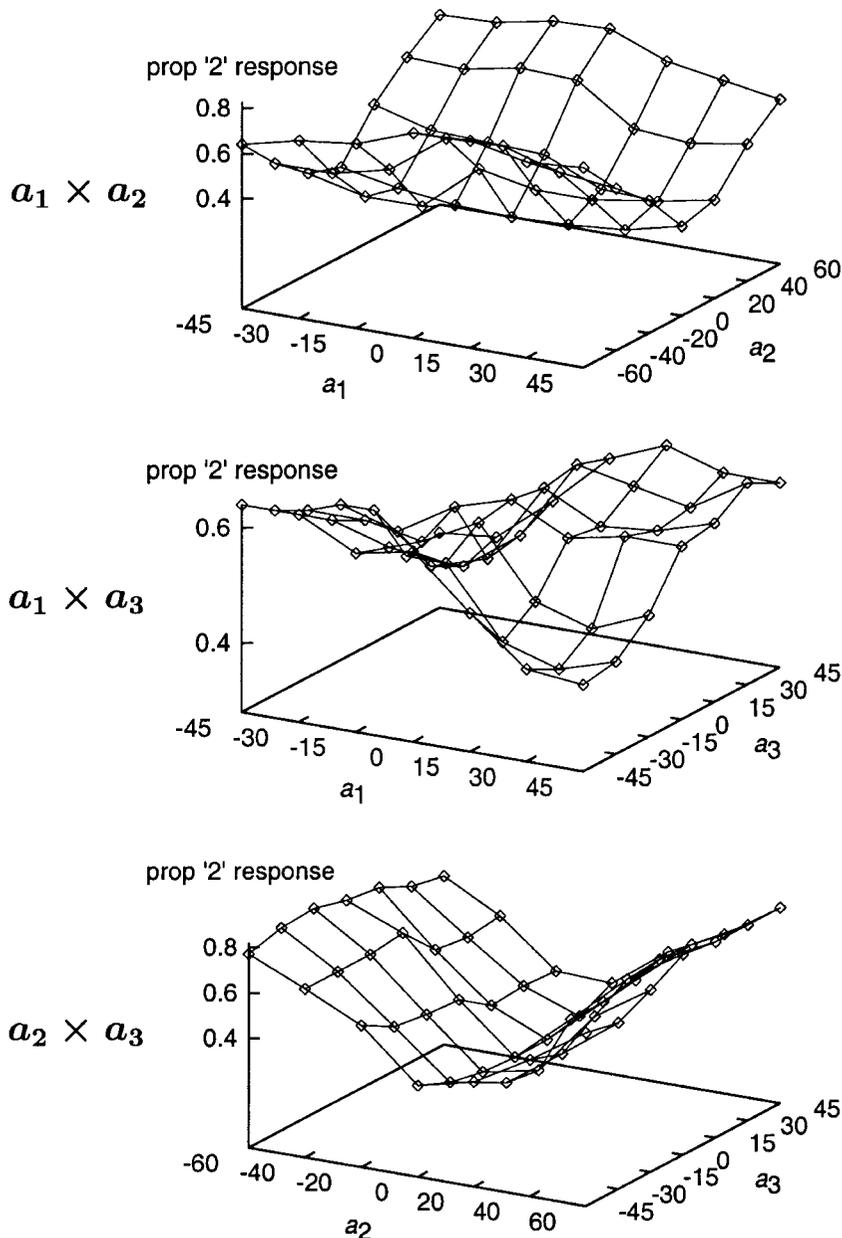


Figure 10. Two-way interactions in the two-contours task: $a_1 \times a_2$ (top), $a_1 \times a_3$ (middle), and $a_2 \times a_3$ (bottom).

case, the use of Gaussian likelihood functions is well motivated by Bayesian theory. The Gaussian distribution is the maximum-entropy distribution given a fixed mean and variance (Bernardo & Smith, 1994). In the present case, this means that, given an assumption that curves tend to continue smoothly, with a certain expectation σ of spread about the expected direction, a Gaussian likelihood function encodes this assumption with the absolute minimum of additional assumptions or information (Jaynes, 1983). This makes the choice of Gaussians a very reasonable one for a visual system in which contours are expected to be curvilinear but in which any more specific knowledge of

the geometry of contours in its environment (as in Brunswik, 1956) is lacking. Ultimately, the close fits reported above suggest that whatever their rationale, Gaussians are the visual system's choice, or very nearly so.

Traditionally, within psychology, perceptual grouping has often been treated as a "subjective" task—that is, one without a definitely correct or incorrect answer. Indeed, the "goodness" of a grouping interpretation has sometimes been treated by Gestalt writers as if it were virtually an aesthetic attribute. Conversely, it is possible to view grouping as the solution to a definite problem—namely, the identification of visual elements in the image

Table 2
Estimated Parameters of the Bayesian Models

Parameter	Corners Task				Parameter	Two-Contours Task			
	Estimate	SE	F(6,637)	p <		Estimate	SE	F(6,343)	p <
σ_3	16.64°	0.72	102.36	.000001	σ_3	13.21°	0.94	24.86	.000001
σ_4	27.59°	0.58	556.24	.000001	σ_4	29.50°	1.24	183.30	.000001
r	0.27	0.02	53.07	.000001	r	0.32	0.03	22.62	.000001
p(H _c)	0.23	0.02	35.93	.000001	p(H ₂)	0.19	0.05	5.12	.000005
p(H _c)	0.06	0.003	382.06	.000001	p(H ₂)	0.18	0.02	92.96	.000001
h	0.79	0.01	282.21	.000001	h	0.77	0.01	20.39	.000001

that arose from the same physical source in the scene—for example, the same object edge or surface. This view is attractive from a computational point of view, in that it provides the beginnings of a rationale for the selection of algorithms: Prefer those algorithms that tend to solve that problem successfully. Clearly, this is the assumption implicit in the Bayesian theory proposed above, in which the observer explicitly chooses between the hypothesis that a given element arises from one curve and the hypothesis that it arises from another. Nevertheless, some care must be taken to avoid circularity here, because the question of which elements *in the scene* are “truly” part of the same object does not bear a completely objective answer and would normally be answered by appealing perceptual grouping (perhaps from another view of the same scene, at a different viewing angle, or at a different time). Still, one can certainly imagine that some grouping interpretations would be more successful than others in terms of their consistency with the later actions of the observer, such as manipulation of the objects observed in the scene.

In the case of Bayesian grouping, this raises a puzzle: If Bayesian theory tends to yield the “correct” answer, what exactly is correct about it in the case of grouping? An argument drawn from the historical debate over Bayesianism provides one possible answer. The Bayesian grouping interpretation is not necessarily an objectively optimal interpretation of the world, but it does represent a provably optimal use of the observer’s prior knowledge and beliefs (Bernardo & Smith, 1994; Jaynes, 1983): in this case, the visual system’s beliefs about the geometric form of smooth curves, as here embodied by the likelihood functions L_3 and L_4 . That is, Bayesian theory yields grouping percepts that parse the world in a way most harmonious with the visual system’s implicit mental model of smooth curves. That this tends to lead to successful percepts reflects both the aptness of the subjective priors and the optimality of Bayesian theory’s use of them.

In this sense, Bayesian grouping theory is not truly normative, because the priors and likelihood functions are motivated by extra-Bayesian considerations (Jepson, Richards, & Knill, 1996; Richards, Jepson, & Feldman, 1996) and, in some surprising respects, even run counter to mathematically correct statistical reasoning (Feldman, 1996). Nevertheless, these functions are apparently manipulated by neural machinery in an almost strictly Bayesian fashion. This realization sheds new light on the

Gestalt principle of *good continuation*, which, along with other Gestalt principles, is conventionally held up as a cornerstone of human perceptual grouping. As has often been pointed out, such principles really describe a particular behavioral tendency, rather than explain or justify it. The link to Bayes’ rule demonstrated here suggests that the principle of good continuation, and, given the theoretically universal range of applicability of Bayesian reasoning, perhaps other Gestalt principles, can indeed be regarded as a rational strategy for perception. The significance of this point, in the context of the recent debate over the evolutionary adaptiveness of human mental strategies and heuristics, is obvious. Indeed, the natural next step is to extend the mathematics presented above to other types of perceptual grouping (e.g., region grouping, amodal completion, etc.), with the hope that the close fit to human judgments found here can be duplicated outside the narrow case of contour integration.

REFERENCES

- ADELSON, E. H. (1993). Perceptual organization and the judgment of brightness. *Science*, **262**, 2042-2044.
- BERNARDO, J. M., & SMITH, A. F. M. (1994). *Bayesian theory*. Chichester, U.K.: Wiley.
- BRUNSWIK, E. (1956). *Perception and the representative design of psychological experiments*. Berkeley: University of California Press.
- BÜLTHOFF, H. H., & YUILLE, A. L. (1991). Bayesian models for seeing shapes and depth. *Comments on Theoretical Biology*, **2**, 283-314.
- CAELLI, T. M., & UMANSKY, J. (1976). Interpolation in the visual system. *Vision Research*, **16**, 1055-1060.
- COMPTON, B. J., & LOGAN, G. D. (1993). Evaluating a computational model of perceptual grouping by proximity. *Perception & Psychophysics*, **53**, 403-421.
- 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. (2000). Bias toward regular form in mental shape spaces. *Journal of Experimental Psychology: Human Perception & Performance*, **26**, 1-14.
- 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.
- GILBERT, C. D. (1995). Dynamic properties of adult visual cortex. In M. S. Gazzaniga (Ed.), *The cognitive neurosciences* (pp. 73-90). Cambridge, MA: MIT Press.
- GILCHRIST, A. L. (1977). Perceived lightness depends on perceived spatial arrangement. *Science*, **195**, 185-187.
- GLASS, L. (1969). Moiré effects from random dots. *Nature*, **223**, 578-580.
- HE, Z. J., & NAKAYAMA, K. (1992). Surfaces versus features in visual search. *Nature*, **359**, 231-233.

- JAYNES, E. T. (1983). E. T. Jaynes: *Papers on probability, statistics and statistical physics* (R. D. Rosenkrantz, Ed.). Dordrecht: D. Reidel.
- JAYNES, E. T. (1993). A backward look to the future. In W. T. Grandy & P. W. Milonni (Eds.), *Physics and probability* (pp. 261-275). Cambridge: Cambridge University Press.
- JEPSON, A., RICHARDS, W. A., & KNILL, D. C. (1996). Modal structure and reliable inference. In D. C. Knill & W. Richards (Eds.), *Perception as Bayesian inference* (pp. 63-92). Cambridge: Cambridge University Press.
- KNILL, D., & RICHARDS, W. (Eds.). (1996). *Perception as Bayesian inference*. Cambridge: Cambridge University Press.
- KUBOVY, M., HÖLCOMBE, A. O., & WAGEMANS, J. (1998). On the lawfulness of grouping by proximity. *Cognitive Psychology*, **35**, 71-98.
- KUBOVY, M., & WAGEMANS, J. (1995). Grouping by proximity and multistability in dot lattices: A quantitative gestalt theory. *Psychological Science*, **6**, 225-234.
- LANDY, M. S., MALONEY, L. T., JOHNSTON, E. B., & YOUNG, M. (1995). Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research*, **35**, 389-412.
- LINK, N. K., & ZUCKER, S. W. (1987). Sensitivity to corners of flow patterns. *Spatial Vision*, **2**, 233-244.
- LIU, Z., KNILL, D. C., & KERSTEN, D. (1995). Object classification for human and ideal observers. *Vision Research*, **35**, 549-568.
- MAMASSIAN, P., & LANDY, M. S. (1998). Observer biases in the 3D interpretation of line drawings. *Vision Research*, **38**, 2817-2832.
- NAKAYAMA, D., & SHIMOJO, S. (1992). Experiencing and perceiving visual surfaces. *Science*, **257**, 1357-1363.
- PARENT, P., & ZUCKER, S. W. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, **11**, 823-839.
- PIZLO, Z., SALACH-GOLYSKA, M., & ROSENFELD, A. (1997). Curve detection in a noisy image. *Vision Research*, **37**, 1217-1241.
- RICHARDS, W. A., JEPSON, A., & FELDMAN, J. (1996). Priors, preferences, and categorical percepts. In D. C. Knill & W. Richards (Eds.), *Perception as Bayesian inference* (pp. 93-122). Cambridge: Cambridge University Press.
- SHIMOJO, S., & NAKAYAMA, K. (1990). Amodal representation of occluded surfaces: Role of invisible stimuli in apparent motion correspondences. *Perception*, **19**, 285-299.
- SMITS, J. T., & VOS, P. G. (1987). The perception of continuous curves in dot stimuli. *Perception*, **16**, 121-131.
- SMITS, J. T., VOS, P. G., & VAN OEFFELEN, M. P. (1985). The perception of a dotted line in noise: A model of good continuation and some experimental results. *Spatial Vision*, **1**, 163-177.
- TREISMAN, A. (1982). Perceptual grouping and attention in visual search for features and for objects. *Journal of Experimental Psychology: Human Perception & Performance*, **8**, 194-214.
- WEISS, Y., & ADELSON, E. H. (1998). *Slow and smooth: A Bayesian theory for the combination of local motion signals in human vision* (A.I. Memo No. 1624). Cambridge, MA: Massachusetts Institute of Technology, Artificial Intelligence Laboratory.
- YUILLE, A. L., & BÜLTHOFF, H. H. (1996). Bayesian decision theory and psychophysics. In D. C. Knill & W. Richards (Eds.), *Perception as Bayesian inference* (pp. 123-162). Cambridge: Cambridge University Press.
- ZUCKER, S. W., STEVENS, K. A., & SANDER, P. (1983). The relation between proximity and brightness similarity in dot patterns. *Perception & Psychophysics*, **34**, 513-522.

(Manuscript received August 15, 2000;
revision accepted for publication January 31, 2001.)