



Curvilinearity, Covariance, and Regularity in Perceptual Groups

JACOB FELDMAN*

Received 9 November 1995; in revised form 17 September 1996; in final form 1 April 1997

A curvilinear pattern among a series of visual items (e.g., dots) can be regarded as a kind of probabilistic inference, in which each consecutive angle, regarded independently, is more nearly collinear than would be expected by chance alone. This paper investigates judgments of curvilinearity as a function of the joint distribution of successive inter-dot angles. Subjects were asked to classify 4- and 5-dot configurations as having been generated by a curvilinear generating process, vs independently. Their results are distributed as a gaussian over the inter-dot angles with mean 0 deg (collinear), with a negative correlation between successive angles, but negligible correlation between non-successive angles. This suggests that curvilinearity is evaluated in a 4-dot window moving along the chain of dots, evaluating collinearity and smoothness but ignoring higher-order relationships. Moreover, the probabilistic model provides a remarkably precise numeric prediction of the magnitude of the correlation. Subjects also showed a reliable preference for equal spacing of dots along the virtual curve. © 1997 Elsevier Science Ltd

Grouping Collinearity Curvature

THE INTERPRETATION OF COLLINEARITY AND CURVILINEARITY

One of the most salient effects in perceptual grouping is that collinear or nearly collinear visual items tend to cohere, for example causing virtual lines and curves to emerge perceptually from a random field of dots. This tendency, called "good continuation" by the Gestaltists, has never been fully understood. It has been linked to orientation-tuned local operators in the visual system (Glass, 1969; Stevens, 1978; Caelli & Julesz, 1978; Prazdny, 1984; Brookes & Stevens, 1991), as well as to larger conglomerations of cells (Dodwell, 1983; Field *et al.*, 1993). Nevertheless, a theoretical account of the collinearity preference, for example, predicting the exact magnitude of the response to various deviations from perfect collinearity, has been elusive.

Theoretical attention has focused on the inference of a virtual curve or trace underlying the discrete dots in the chain (Zucker & Davis, 1988), especially in computer vision, where such an inference is a necessary component of the integration of contours from sparse image data (Zucker, 1985; Parent & Zucker, 1989). Collinear or nearly collinear dots can be regarded as a proximate cue to the presence of a distal generating curve. In this view a collinear or curvilinear pattern amounts to a kind of "regularity," a special configuration to which the visual system is especially attuned, not unlike mirror symmetry

(Barlow & Reeves, 1979), skewed symmetry (Wagemans, 1993), and other higher-order patterns (Wagemans *et al.*, 1993).

A probabilistic approach

To fully understand the interpretation of collinearity, one would like to be able to precisely quantify the relationship between proximate collinearity and the presence of a distal curve—that is, how straight is straight "enough" to justify inferring a curve? When viewed as a Bayesian probabilistic problem, a necessary component of the answer is the prior distribution of dot configurations that obtains when a curve is actually present—or, at least, the observer's subjective expectations about this distribution. This distribution was investigated for the minimal case of three dots in Feldman (1993) and Feldman (1996). Of course, subjects have no conscious knowledge of these distributions; rather, the idea is that the probabilistic machinery captures mathematically the tacit expectations that are physically realized in the relevant neural hardware. Despite this added layer of abstraction, the probabilistic approach provides an extreme numerically precise account of observers' collinearity judgments in the 3-dot case. The current study extends this analysis to four and more dots.

To illustrate the approach, consider a group of dots in a curvilinear pattern (Fig. 1). Because each individual dot conveys no directionality, the curvilinearity is only an interpretation placed on the dots' relative positions, and does not correspond in any simple way to physical

*Department of Psychology, Center for Cognitive Science, Rutgers University, New Brunswick, NJ 08903, U.S.A. [Fax: +1-908-445-0634; Email: jacob@ruccs.rutgers.edu].




FIGURE 1. A group of dots arranged in an apparently curvilinear pattern, and the resulting series of inter-dot angles a_1 , a_2 and a_3 .

structure present in the image. Rather, the curvilinearity resides in a series of apparent coincidences: the angles between successive dots, a_1 , a_2 , and a_3 , are each nearer to collinear than one would expect "by chance alone." More specifically, if these dots were all generated independently, then this series of nearly collinear angles would have to be regarded as a highly suspicious coincidence; whereas if they were generated by some kind of organized curvilinear process, then the observed pattern is about what one would expect. The logic is similar to that of "non-accidental features" (Binford, 1981; Lowe, 1987). From a probabilistic standpoint, each potentially independent event (here, the angle formed by each contiguous group of three dots) would be regarded as a single trial, with angle drawn from some distribution; while the overall pattern would be regarded as a series of

such trials (with n dots, $n-2$ trials). Foster (1983) has shown that human perception of curvilinear segments exhibits categorical effects, even when the underlying stimulus space is carefully arranged to be metrically uniform. In the current case of discrete dots, the "curvilinear" category can be thought of as, in effect, a probability distribution defined over the series of inter-dot angles (see Ashby & Perrin, 1988 and Nosofsky, 1991). There are then two straightforward research questions concerning dot patterns that are perceived as curvilinear:

1. What is the expected distribution of each individual angle (three dots)?
2. What is the joint distribution of multiple angles ($n > 3$ dots)?

The first question, the distribution of judgments in the minimal 3-dot case, was addressed in Feldman (1993) and Feldman (1996). Subjects were asked whether three dots were generated by a curvilinear generating process, as opposed to having been generated independently (translated into suitably comprehensible terms; see Procedure, below). The probability of a "yes" (curvilinear) response as a function of angle can be taken to represent their model of the likelihood function due to a curvilinear dot-generating process—that is, the expectation of each successive angle in the curvilinear sequence of Fig. 1. The distribution turns out to be approximately gaussian,* centered at 0 deg (straight†), with a standard deviation of about 53 deg, tailing off to minimum at about 120 deg (Fig. 2). This shape makes sense, by the following argument. The 0 deg case represents the modal outcome from a curvilinear process, in which the inference to a curvilinear process is maximally justified.

*More precisely, the distribution turns out to be the most-nearly gaussian among the family of posterior functions with gaussian priors. This function can be closely approximated as a gaussian with a free height term. In the context of a maximum-likelihood interpretation, this function then plays the role of both prior and posterior. Hence, in what follows we ignore this distinction, and assume that judgment distributions can be regarded as likelihood.

†All inter-dot angles in this paper are measured from straight, as in Fig. 1.




FIGURE 2. Probability of a "curvilinear" response as a function of angle, in the 3-dot case (Feldman, 1996).

On the other hand, the 120 deg case (i.e., an equilateral triangle), because of its symmetry, represents the completely ambiguous case where the order in which the three dots were generated is completely unrecoverable, and hence evidence for the curvilinear interpretation is maximally weak. Thus, 0 and 120 deg are maximally distant in the underlying space of hypotheses (in a sense that does not depend on the choice of metric). Note, however, that the mean judgment was not unity at 0 deg and was not zero at 120 deg, reflecting some residual uncertainty about the classification of even these most extreme cases.

In this account, subjects treated 120 deg as the prototype for generic triplets. This is remarkable because 120 deg is not, in fact, the modal case when three dots are generated randomly and independently [Kendall & Kendall, 1980; see Feldman (1996) for discussion], but rather is quite unlikely even then. In effect, subjects sought to create a maximally extreme contrast between two hypotheses: regularity (collinearity) and genericity (non-regularity or independence). To achieve this, they confabulated a prior distribution for the independent hypothesis based on the idea of maximum distance from a collinear configuration—maximum genericity. In the end, they can be regarded as performing Bayesian inference only by appeal to some extra-Bayesian principles: (a) for the curvilinear prototype, the angle distribution is centered on the case for which the inference or regularity is most justified; while (b) for the “random” (generic) prototype, the angle distribution is centered on the case in which the inference of regularity is least justified. These “regularity principles” are worth emphasizing, because, as will be seen below, they can be used to derive numeric predictions for the case of $n > 3$.

This paper addresses the second question: with more than three dots, and hence more than one angle in the sequence, what is the joint distribution of successive angles? The key idea intuition here is that in a coherent, curvilinear process such as an edge, one would not expect successive angles to be independent. Rather, within a genuine curvilinear pattern, successive near-collinear angles tend to follow one after another, suggesting some kind of predictive power from one to the next—i.e., some degree of covariance among successive angles. By contrast, in an unstructured cluster of dots, one would expect the various angles to be independent, by definition. It is often remarked that perceptual hypotheses of “structure” depend in some way on covariation among distal variable in the environment. Here is the potential to establish a concrete link between one such hypothesis, that of collinearity, and the literal mathematical notion of covariance.

The critical experiment is exactly like the 3-dot experiment, except with more dots. Results for 4- and 5-dot cases are reported below. Again, the task is to judge whether the given dots were generated by a curvilinear process, or independently.

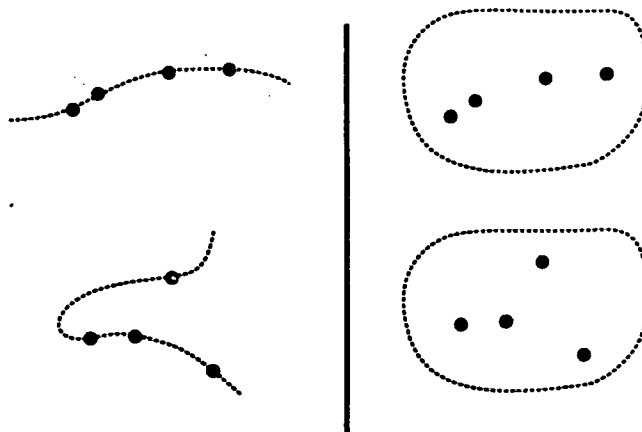


FIGURE 3. Examples of worms (left) and flatfish (right) shown to subjects (4-dot case).

EXPERIMENT

Method

Subjects. Subjects were naïve paid volunteers drawn from the university community, mostly undergraduate students from a variety of disciplines. There were 20 subjects in the 4-dot experiment, and 16 (different) subjects in the 5-dot experiment.

Procedure. Subjects were instructed that they were to be presented with configurations of four (five) dots, to be interpreted as markings on two equinumerous species of bottom-dwelling ocean creatures: a kind of worm, whose four (five) markings were randomly placed along its curvilinear body; and an approximately round flatfish, whose four (five) marking were randomly placed on its surface (Fig. 3). The instructions explicitly drew subjects' attention to the fact that any configuration could be either type of creature: worms can curl up arbitrarily, while the dot markings on flatfish might just “happen” to fall nearly in a line. As illustration, subjects were shown two worms and two flatfishes, arranged so that exactly the same two-dot configurations were given as an example of both categories (see Fig. 3).

Thus, subjects understood that their task was to rate each dot configuration for how likely it was to have been generated by by a curvilinear process. On each trial, subjects responded by choosing a number from 1 to 5, meaning (1) “almost definitely a flatfish”; (2) “probably a flatfish”; (3) “either one, with about equal probability”; (4) “probably a worm”; and (5) “almost definitely a worm”. The experimental session was preceded by eight practice trials.

Design. There are two variables for each configuration: angle profile and inter-dot distance profile. For the 4-dot case the angle profile consists of two angles, a_1 and a_2 , and three inter-dot distances L_1 , L_2 , and L_3 [Fig. 4(a)]. For the 5-dot case the angle profile consists of three angles, a_1 , a_2 and a_3 , and four inter-dot distances, L_1 , L_2 , L_3 , and L_4 .

4-dot case. There were 13 levels of angle (from -90 to 90 deg in increments of 15 deg) in each of a_1 and a_2 ,




FIGURE 4. Parameters for (a) 4-dot stimuli; (b) 5-dot stimuli.


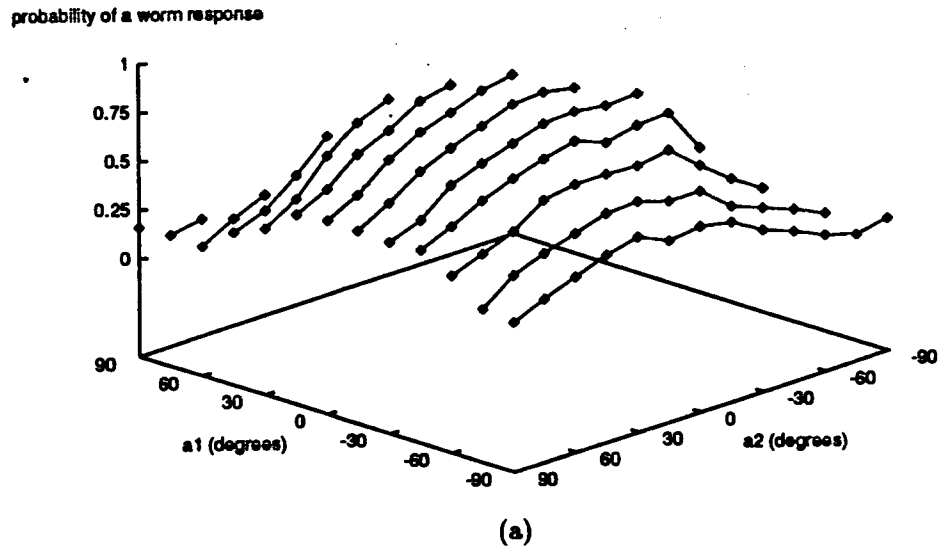


FIGURE 5. Results for $n = 4$ plotted by angle: (a) surface plot; (b) contour plot. Here the data have been reflected about the $a_1 = a_2$ axis to facilitate the contour plot. Contour levels shown range from 0.125 through 0.875 (inclusive) in increments of 0.125 (corresponding to one-half unit in the subjects' response).

using each possible pair only once, for a total of $91 (= 13 + (13)(12/2))$ distinct types of angle profile (angle by angle). To keep the total number of trials down to a reasonable level, only five types of inter-dot distance levels were used: same-same-same (i.e.,

$L_1 = L_2 = L_3$), different-same-same (i.e., $L_1 \neq L_2 = L_3$, etc.), same-different-same, same-same-different, and different-different-different. These five types can be classified by how many independent equal relationships they include: 0 (*ddd*), 1 (*dss*, *sds*, and *ssd*), and





FIGURE 6. Theoretical model of 4-dot angle results: (a) surface plot; (b) contour plot. Contour levels are as in Fig. 5.

TABLE 1. Estimated parameters of the model fitted to the 4-dot angle data

Parameter	Estimate	Standard error	95% confidence interval
<i>h</i>	0.9415	0.0099	0.9218-0.9611
<i>s</i>	74.80	0.9452	72.92-76.67
<i>r</i>	-0.2634	0.01783	-0.2988--0.2280

2 (sss).* "Same" inter-dot distances were all a standard length (80 pixels, subtending 3.68 deg of visual angle), while "different" inter-dot distances were different from

*Note that though sss includes three sames, only two of them are independent.

standard by some non-zero integral power of 1.2, chosen randomly from the interval [-4, 4].

The angle scheme was crossed with the length-profile scheme for a total of 455 (= 91 × 5) trials. Each subject saw all trials.

5-dot case. There were 7 levels of angle (from -72 to 72 deg in increments of 24 deg) in each of *a*₁, *a*₂ and *a*₃. All distinct pairs of *a*₁ and *a*₂ times all levels of *a*₃ were used, for a total of 196 [= (7 + (7)(7/2)) × 7] distinct types of angle profile (angle by angle by angle). Because of this large number of angle levels, only two types of inter-dot length profile were used: same and different. As above, "same" inter-dot distances were all a standard length (80 pixels, 3.68 deg of visual angle), while "different" inter-dot distances were different from standard by some




FIGURE 7. Frequency histograms of individual subjects' values of (a) h (b) s (c) r and (d) R^2 (4-dot case).

non-zero integral power of 1.2, drawn randomly from $[-3, 3]$. (The smaller range was in order to limit the maximum size of the objects.)

The angle scheme was crossed with the length profile scheme for a total of 392 ($= 196 \times 2$) trials. Each subject saw all trials.

Viewing distance was 35 cm. Dots were black circular patches 0.23 deg wide on a white screen with high contrast. On each trial, the configuration was rotated in the plane by a random angle. Subjects were free to take as long as they wanted to respond, and were allowed to move their eyes freely.

RESULTS

The results for angle and the results for inter-dot distance will be discussed separately.

Angle results

4-dot case. The results for the 4-dot case are shown in Fig. 5, plotted by a_1 and a_2 , collapsing over all inter-dot distances. The vertical axis is subjects' average response, normalized to the interval (0, 1) (i.e., converted to a probability judgment). Hence, the plot can be regarded as the probability distribution of "worms" as a function

of the two angles: that is, the joint distribution of a_1 and a_2 .

The distribution was modeled as a bivariate gaussian with a free height parameter and equal standard deviations:

$$p(a_1, a_2) = h \exp\left(-\frac{1}{2} \mathbf{a}' \Sigma^{-1} \mathbf{a}\right) \quad (1)$$

in which Σ is the 2-by-2 covariance matrix, \mathbf{a} is the vector of angles

$$\mathbf{a} = [a_1 \ a_2], \quad (2)$$

and \mathbf{a}' is its transpose. This model assumes that each a_i has an expected mean of 0 deg, as one would expect for the "worm" distribution. This model has 3 deg of freedom: the free height h , the single standard deviation s , and the correlation coefficient r . The model was fit to the data by Levenburg-Marquardt. Estimates, standard errors, and 95% confidence intervals for each of the parameters are given in Table 1. The estimated model, shown in Fig. 6, fit the data extremely well, $F(3,87) = 912.169$, $P < 0.0001$, $R^2 = 0.9692$.* The sub-

*Note that this R^2 , a measure of fit between the data and the model, is unrelated to r , which here is a parameter of the model.

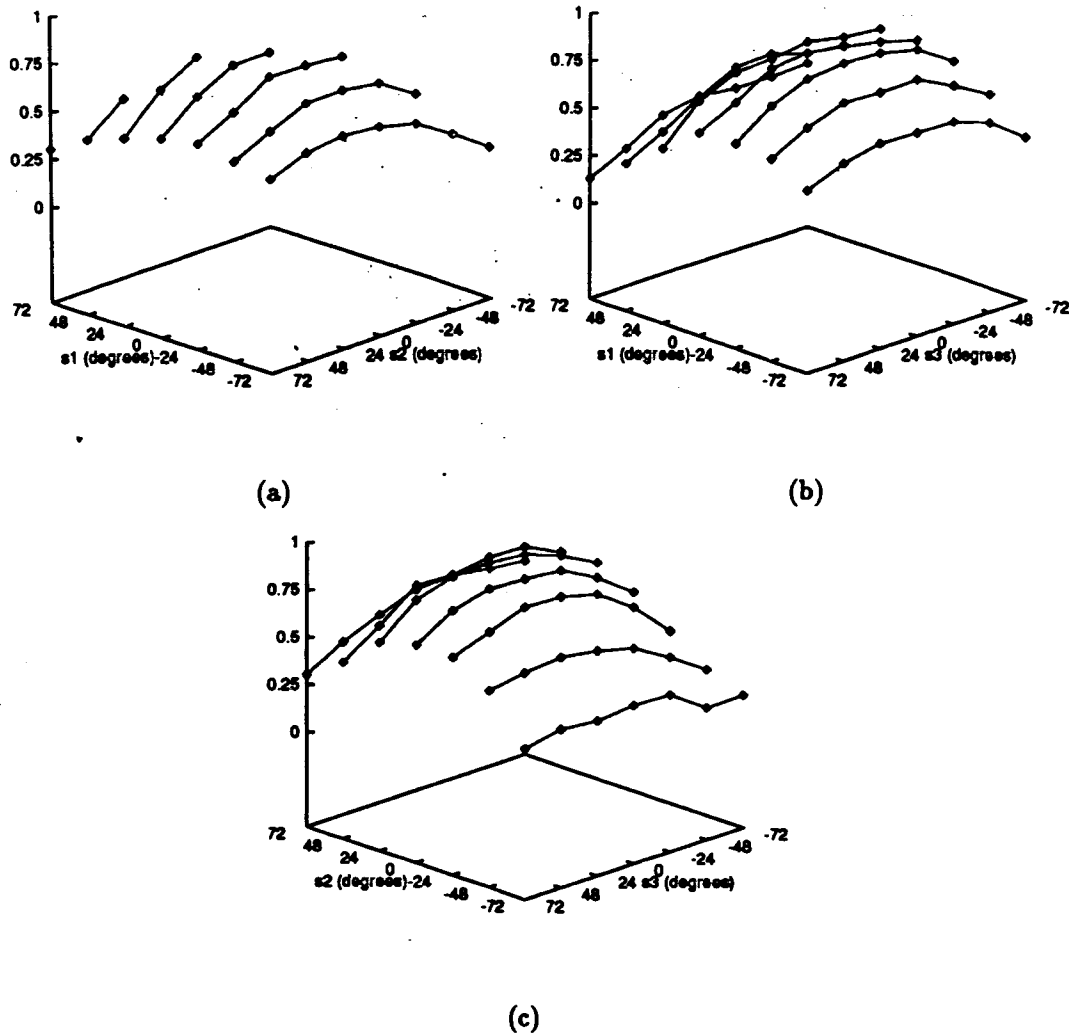


FIGURE 8. Results for $n = 5$ plotted by angle: (a) by a_1, a_2 (b) by a_1, a_3 (c) by a_2, a_3 .

jects' mental "worm" distribution very closely approximates a gaussian centered on $a_1 = a_2 = 0$ deg (fits to individual subjects' data are discussed below). This result is quite striking when one considers the total freedom subjects had to invent arbitrary decision criteria for what is, after all, a completely subjective task.

The significant correlation ($r = -0.2634$) means that subjects effectively treated collinearity as a covariance between successive angles. Intriguingly, the correlation is negative; this will be explained by a model presented below.

Variation among subjects. The gaussian model can be fitted to each individual subject's data as well, yielding individual estimates of the model parameters. Figure 7 shows frequency histograms for each of the parameters h , r , and s , as well as for the measure of fit R^2 . For each of the three parameters, the plot shows a fairly narrow spread about the value fitted overall. In addition, the

range of fits to the model is quite good, with most at about $R^2 = 0.8$ or above. This confirms that the overall fit of the gaussian model was in fact due mostly to individual subjects fitting it, with only a small amount of the fit (about 10%) being due to regression to a gaussian stemming from the Central Limit Theorem.*

5-dot case. In the 5-dot case, there are three angles a_1 , a_2 , and a_3 , so their joint distribution cannot be plotted in three dimensions. Instead, each of the three pairwise joint distributions is plotted (Fig. 8). The three-way joint distribution was modeled as a trivariate gaussian:

$$p(a_1, a_2, a_3) = h \exp\left(-\frac{1}{2} \mathbf{a}' \Sigma^{-1} \mathbf{a}\right) \quad (3)$$

in which \mathbf{a} is now the three-component vector of angles, and again Σ is the covariance matrix with equal standard deviations. This model has five degrees of freedom: the free height h , the standard deviation s , and the three pairwise correlation coefficients r_{12} , r_{23} , and r_{13} . This model again fit the data extremely well, $F(5, 190) = 577.53$, $P < 0.0001$, $R^2 = 0.9383$. Estimates, standard errors, and 95% confidence intervals for each of these parameters are given in Table 2. Note that the

*Notice that individual R^2 's, unlike individual model parameters, do not need to be clustered about the overall value for all subjects. Instead, the overall fit will tend to regress towards a gaussian, inevitably yielding lower R^2 's for individual subjects.

TABLE 2. Estimated parameters of the model fitted to the 5-dot data

Parameter	Estimate	Standard error	95% confidence interval
h	0.9586	0.0091	0.9406–0.9766
s	91.90	1.360	89.213–94.579
r_{12}	-0.2955	0.0195	-0.3340–-0.2571
r_{23}	-0.3167	0.0200	-0.3561–-0.2773
r_{13}	-0.0613	0.0231	-0.1069–-0.0157


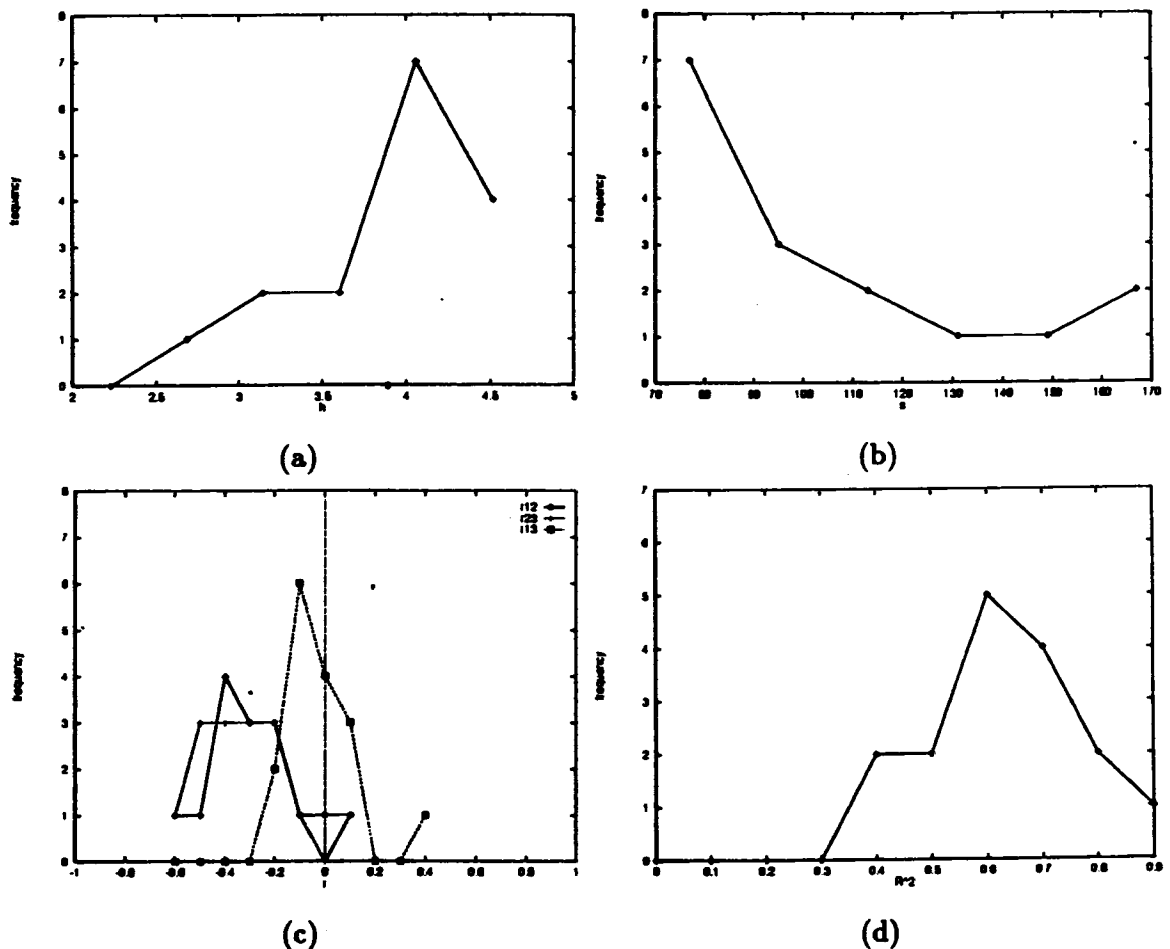


FIGURE 9. Curvilinearity is evaluated within a moving 4-dot window.

Here, as in the 4-dot case, subjects modeled the worm distribution as a gaussian centered on $a_1 = a_2 = a_3 = 0$. Again curvilinearity took the form of a correlation between successive angles ($r_{12} = -0.2955$, $r_{23} = -0.3167$). Intriguingly, the correlation between non-successive angles ($r_{13} = -0.0613$) is significantly lower (though still significantly different from 0 at the 0.05 level, and just barely different from 0 at the 0.01 level). This suggests that human observers classify a chain of dots as curvilinear if (a) each individual angle in the chain is near enough to collinear; (b) each pair of adjacent angles is near enough to curvilinear. This means curvilinearity is evaluated using a 4-dot window, while any higher-order relationships are (almost completely) ignored* (Fig. 9). In other words, the sequence of dots is evaluated for local† straightness and local smoothness—the vanishing of the first and second derivatives of the

FIGURE 10. Frequency histograms of individual subjects' values of (a) h ; (b) s ; (c) r_{12} , r_{23} , r_{13} ; and (d) R^2 (5-dot case).

standard deviations are generally somewhat higher than those in the 4-dot case, which in turn were higher than the original 3-dot case. This trend presumably reflects the strengthening of the curvilinear interpretation (entailing a wider curvilinear distribution) warranted by increasing numbers of dots.

tangent vector—but higher-order derivatives are disregarded.

*At least any higher-order *local* relationships: see Kovacs & Julesz (1993) for evidence concerning a global factor, closure.

†“Local” here means in a neighborhood along along the dot chain, not at a single dot.




FIGURE 11. Results for $n = 4$ plotted by inter-dot distances. The abscissa measures degree of regularity in the inter-dot distances: 0 = *ddd*, 1 = (*dss*, *sds*, *ssd*), and 2 = *sss*.

This finding is particularly intriguing in the light of the finding by Wagemans *et al.* (1993) that 4-dot groups play a central role in the detection of symmetry and other kinds of regularity. In their studies, subjects were better able to detect various types of "regular structure" in dot configurations—mirror symmetry, translational symmetry, and rotational symmetry—when the patterns contained structure involving more than two dots at a time. Pair-wise regularities among dots by themselves were not sufficient to allow observers to detect the structure effectively. In particular, Wagemans *et al.* argue that sets of four dots—defining a "virtual quadrilateral"—suffice to allow the observer to bootstrap the impression of regularity rapidly across the image. The current findings have a very similar flavor. A chain of dots is regarded not as a concatenation of simple pair-wise collinearities, but rather is analyzed in overlapping groups of four dots, allowing an accurate assessment of "curve goodness" to propagate rapidly along the dot chain (cf. Field *et al.*, 1993). The use of more than two dots at a time immediately suggests that the visual system is calculating not only local tangents but local curvature as well, an idea that will play a central role in the model presented below.

Variation among subjects. Figure 10 shows frequency histograms of the individual 5-dot subjects' parameters, as well as of the fits (R^2 's) to the gaussian model. Again, most of the individual subjects fit the model well, though not as tightly as in the 4-dot case, with most R^2 's at 0.6 or better. As in the 4-dot case, each of the parameters shows a fairly tight distribution about the overall estimate (somewhat broader for s). Particularly striking is the case of the three correlations (lower left). The two successive correlations, r_{12} and r_{23} , are each clustered around the overall estimates of about -0.3 in each case; while the

non-successive correlation r_{13} is clumped near zero. This plot makes a strong case that the moving 4-dot window described above is actually a universal strategy among the subjects.

Inter-dot distance results. Figure 11 shows inter-dot distance results for the 4-dot case. Probability of a worm response increases monotonically with the number of independent "same" lengths, i.e., with degree of regularity ($F(2,38) = 4.112$, $P = 0.024$). On the other hand, as expected, there was no significant variation within the "1" group ($dss = 0.5545$, $ssd = 0.5511$, $sds = 0.5508$), $F(2,38) < 1$.

Figure 12 shows the corresponding effect in the 5-dot case. Here, the inter-dot length profiles were simply "same" vs "different". Again, "same" configurations were judged more likely to be a worm $F(1,15) = 6.613$, $P = 0.021$.

In both cases, note that the effect of inter-dot length regularity is very small in magnitude compared with the corresponding angle effect. Nevertheless, it is statistically reliable, and goes in the expected direction: equality of inter-dot distances is interpreted as evidence of generation by a curvilinear process. That is, regular behavior of image items is taken to imply common origins. This effect is quite telling, because unlike the corresponding angle effect, it appears to have no rational basis—"worm" configurations cannot, in general, be expected to have more symmetric inter-dot distances than "flatfish" configurations. Rather, the effects signal a purely heuristic preference for regularity.

Model

The negative correlation between successive angles (approximately the same value in the 4-dot case as in




FIGURE 12. Results for $n = 5$ plotted by inter-dot distances.

each of the two successive angle pairs in the 5-dot case) is crucial. In a sense, this one number captures everything there is to say about the human probabilistic model of collinearity. Hence, we would like to be able to explain why it takes on the particular value it does. In particular, why is it negative?

An answer to this question, including an accurate numerical retrodiction of the magnitude of the correlation, is provided by a generalization of the argument used in Feldman (1993, 1996) to account for the corresponding question in the 3-dots case (namely, what is the standard deviation of the angle distribution?). Motivating the argument requires delving in some depth into the geometry of the 4- and 5-dot situations.

What are the regularities? As discussed above, judgments in the 3-dot case revolved about the question of what dot configurations are most and least prototypical of the curvilinear generating process: the former being a "regularity" and the latter a "generic" or non-regular configuration (Jepson & Richards, 1992; Richards *et al.*, 1996; Feldman, 1997a, b). To account for the value of the correlation, we need to generalize this idea to the 4-dot case. (Notice that by the argument given above, we do not then need to move on to the 5-dot case, because observers treat it as two adjoined 4-dot cases.) Clearly, $a_1 = a_2 = 0$ deg is the regularity, the modal form for a prototypical 4-dot worm, and hence is the maximum in subjects' distribution. But the argument also requires knowledge of the maximally non-regular case. Where is the maximally generic point in $\langle a_1, a_2 \rangle$ -space?

Consider a reparameterization of $\langle a_1, a_2 \rangle$ -space, a simple change of variables that rotates the coordinate system by 45 deg (Fig. 13). The new parameters, $a_1 + a_2$ and $a_1 - a_2$, reflect the relationship between the two angles, rather than the individual angles themselves.

Thinking of the dots as having been generated by a curve, each individual angle corresponds to (an approximation of) curvature at the point. Curvature is well known to be perceptually important (Attneave, 1954; Koenderink & Richards, 1988) and to exhibit categorical effects (Foster, 1983). Small differences in curvature are very sensitively computed by human observers (Watt & Andrews, 1982; Wilson, 1985; Foster *et al.*, 1993), even manifesting hyperacuity effects (Fahle, 1991). In the new




FIGURE 13. The new coordinate system, $\langle a_1 + a_2, a_1 - a_2 \rangle$. Also shown is an ellipse with principal axes α and β , which is an isoprobable contour for a bivariate gaussian.

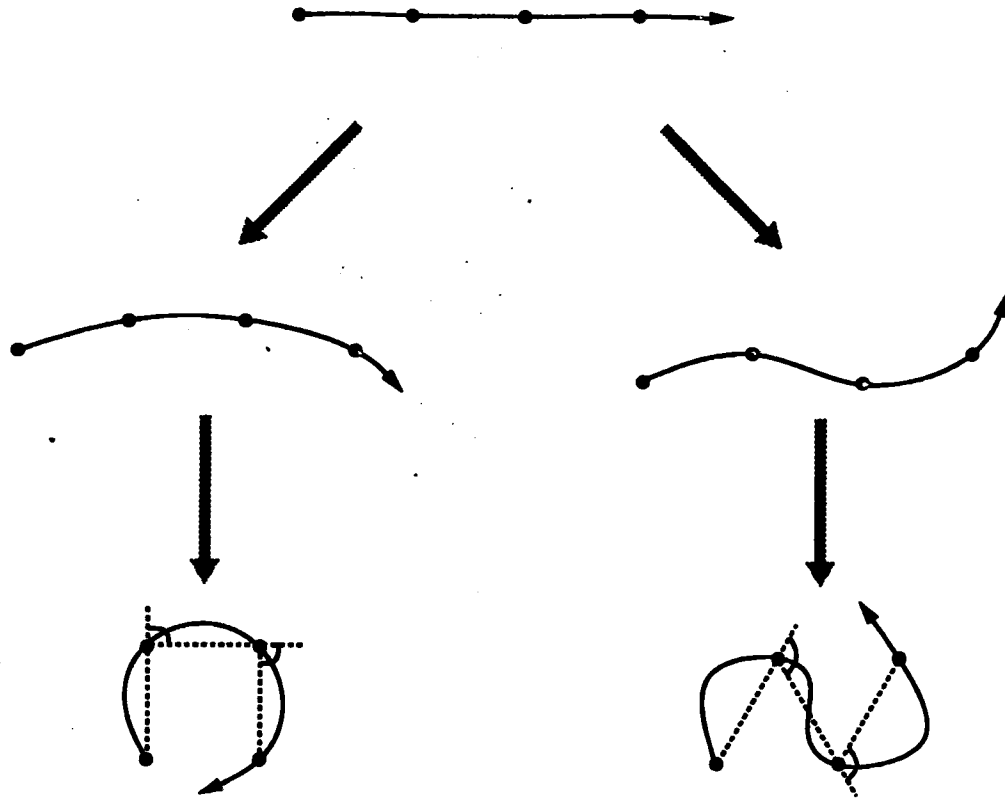


FIGURE 14. The reasoning behind the values of k_{\max} and Δk_{\max} given in the text. A maximally regular prototype "worm" (top) is deformed in the k_{\max} direction (left) and in the Δk_{\max} direction (right) until it reaches an ambiguous configuration (bottom).

parameterization, it is easy to see that $a_1 + a_2$ corresponds to the total curvature k (i.e., total excursion of the tangent direction) and $a_1 - a_2$ corresponds to change in curvature Δk :

$$\begin{aligned} k &= a_1 + a_2, \\ \Delta k &= a_1 - a_2. \end{aligned} \tag{4}$$

What is the modal form for a "worm," i.e., the configuration optimally indicative of having been generated by a curve, expressed in this new coordinate system? It seems quite reasonable to suppose that the answer would be zero total curvature, zero change in curvature, i.e., the worm that bends as little and as consistently as possible. This gives:

$$\begin{aligned} a_1 + a_2 &= 0^\circ \\ a_1 - a_2 &= 0^\circ, \end{aligned} \tag{5}$$

hence $a_1 = a_2 = 0$ deg, our original origin. But unlike the original coordinate system, the new system makes a prediction about the least regular configurations as well. Recall that in the 3-dot case, this turned out to be the equilateral triangle, apparently because in this configuration the curve has bent so much that, because of the symmetry, it is no longer even possible to recover the sequence in which the dots were generated.

Now, with a bit of mathematics, the same argument provides a numeric prediction about the value of the correlation r in the 4-dot case. To see how, consider the

contour plot of a bivariate gaussian [e.g., Fig. 6(b), also shown in Fig. 13]. The isoprobable contour is an ellipse, whose principal axes α and β lie along the eigenvectors of the covariance matrix, which point exactly in the directions of our new parameters $a_1 + a_2$ and $a_1 - a_2$ (Fig. 13). Hence, the reparameterization is a particularly convenient one, though nothing in the following analysis depends on the choice. As in the 3-dot case, the distribution degrades in each of these directions until it reaches a minimum out at the maximally generic point. Hence, we need to identify just how far in the $a_1 + a_2$ direction and in the $a_1 - a_2$ direction a configuration must deform before it becomes completely generic. Call these two extreme points k_{\max} and Δk_{\max} , respectively. It is clear from consideration of Fig. 13 that the ratio of these two distances is precisely the ratio of the principal axes of the ellipse, (i.e., its aspect ratio):

$$\frac{k_{\max}}{\Delta k_{\max}} = \frac{\alpha}{\beta}. \tag{6}$$

This ratio completely determines the shape of the gaussian, and thus must depend on its one free shape parameter, the correlation r . In particular:

$$\frac{\alpha}{\beta} = \sqrt{\frac{1+r}{1-r}}. \tag{7}$$

(A derivation of this fact can be found in the Appendix) Solving for r in terms of k_{\max} and Δk_{\max} , we have

$$r = \frac{(k_{\max})^2 - (\Delta k_{\max})^2}{(k_{\max})^2 + (\Delta k_{\max})^2} \quad (8)$$

Hence in order to predict the expected correlation, it remains only to identify k_{\max} and Δk_{\max} numerically.

As in the 3-dot case, these values are determined by considerations of symmetry. For k_{\max} , the question is: as one bends a straight configuration (Fig. 14, top) in the direction of increased total curvature (Fig. 14, left), at what point is the bend so great that the virtual curve can no longer be identified unambiguously? As can be seen in the figure, the answer is a square, with total bend 180 deg. At this point, the second dot and the fourth dot are both equidistant from the starting dot. For Δk_{\max} , the question is the same except one bends in the direction of increased change in curvature (Fig. 14, right). Here the answer is a equilateral rhombus (i.e., adjoined equilateral triangles), where the total change in curvature is 240 deg.

Inserting the values $k_{\max} = 180$ deg and $\Delta k_{\max} = 240$ deg into Eq. (8) gives a predicted value of:

$$r = -0.28. \quad (9)$$

This value is very close to the empirical value (-0.2634) of r from the 4-dot case, as well to the contiguous correlations r_{12} and r_{23} from the 5-dot case (-0.2955 and -0.3167 , respectively). The average of r , r_{12} , and r_{23} is -0.2919 , only about 0.01 off from the estimate.

It should be noted that subjects' performance cannot be attributed solely to learning the range of the stimuli presented. One can imagine that subjects might learn over the course of the experiment the extremes of the stimulus set, and use these as estimates of k_{\max} and Δk_{\max} . However, the inter-dot angles in the 4-dot case ranged from -90 to 90 deg, yielding a $k_{\max} = \Delta k_{\max} = 180$ deg, entailing $r = 0$ by Eq. (8). Similarly, in the 5-dot case, the angles ranged from -72 to 72 deg, yielding $k_{\max} = \Delta k_{\max} = 144$ deg, and again $r = 0$. Hence, the significant non-zero r s in both cases rule out any explanation stemming solely from the observed stimuli.

In summary, the human model of curvilinearity can be regarded as an expectation of correlation between successive angles in a series of visual items. The strength of the expected correlation is not based on environmental factors, about which the observer has no independent knowledge. Rather it is based on the geometry of how a prototype distorts, which constrains the way the observed configuration is mapped into "worm space". The configuration's location in this space is always defined relative to two opposite poles: the ideally regular form—the worm prototype—and the ideally non-regular form—the generic configuration.

It should be understood that subjects presumably have no conscious knowledge of these complex geometric and probabilistic considerations. Rather the mathematical machinery is an attempt to explain the precise character and magnitude of human observers' unconscious judg-

ments. These very computations might well be realized in some simpler neural substrate—just as in edge detection, where computations that can only be understood completely in the language of calculus are actually carried out using simple combinations of receptive fields. In particular, the neural computation of local curvature has been attributed to endstopped cells (Dobbins *et al.*, 1988; Dobbins *et al.*, 1989); the k and Δk parameters in the current model could easily correspond to excitatory summation and lateral inhibition, respectively, of such cells. The model, complex though its derivation may be, simply describes how such a simple neural computation might correspond to the perceived degree of coherence or regularity of a dot chain.

Hence, while the principal thrusts of the calculation are explained at this "competence" level, it is indeed quite possible that certain aspects of human performance actually depend on details of the physical or algorithmic realization. For example, the geometrical arguments above are by their nature scale-invariant. Indeed, scale invariance might well be the "ideal" behavior of the system. In fact, Takeichi (1995) has recently presented evidence that under certain circumstances human computation of curvature may, in fact, be approximately invariant to scale. Yet ultimately so much of the visual system is modulated by scale that perfect scale invariance seems unlikely. Only future experiments can actually answer this question, but if responses did vary with scale this would presumably simply reflect some unknown details of the way the competence theory was implemented in the visual system.

SUMMARY

Taken together, the angle and distance results paint a very specific picture of the categorization of dot configurations. When classifying dot clusters with regard to curvilinearity, human observers act as if they move a 4-dot window along the dot series, evaluating regularity within the window (Fig. 9). Here "regularity" means

- (a) Curvature vanishes
- (b) Change in curvature vanishes
- (c) Difference in inter-dot distances vanishes.

Components (a) and (b) can be thought of as the vanishing of the Mahalanobis distance (distance scaled by variance) in $\langle a_1, a_2 \rangle$ -space between the configuration and the origin, which is the "worm" prototype. This suggests an empirically motivated representation for dot configurations, in which each successive 4-dot window is mapped to a single two-component vector in the variance-scaled space. It should be possible to design a simple filter to compute this mapping. This suggests a strong empirical test of the theory: judged similarity between 4-dot configurations ought to be a monotonic function of distance in this space. This question will be the topic of future research. Formally, mapping configurations into this 2-D space is exactly equivalent to

evaluating the correlation between the inter-dot angles. In this limited domain, then, "regularity" corresponds literally to covariance.

CONCLUSION

Witkin & Tenenbaum (1983) have argued that perceptual grouping subserves the recovery of "structure" from the image—coordinated behavior on the part of potentially independent visual items, which tends to imply common origins. It would make sense for this inference to be a Bayesian one, if only the observer had access to accurate probabilistic information about the priors: how often does each type of structure occur, what is the probability distribution on each type's outputs, and so forth. As is commonly pointed out, this approach suffers from the fact that this information is often unavailable to the observer. Historically, Bayesians have often tried to supply the missing information via some version of the "Principle of Indifference"—two hypotheses about which one knows nothing are by default assigned equal probabilities. This principle has often been regarded as arbitrary, and the resulting "subjective probabilities" unsound. In its best applications, though, the Principle is really a kind of symmetry principle. A coin whose two sides are physically indistinguishable has, by virtue of the symmetry, equal probabilities of heads and tails. This conclusion seems both justified and empirically correct.

The study reported in this paper suggests that the recovery of curvilinearity, viewed as a problem of probabilistic inference, has a similar (albeit more complex) form. Extreme points for the distributions—crucial for establishing the exact magnitude of the curvilinearity judgment—were derived from symmetry arguments. In a sense, the observer has no choice but to depend on some sort of idealized argument in order to construct priors, because the true environmental probabilities—e.g. the true probability of a given chain of dots having arisen from a curvilinear process—are unavailable, vary arbitrarily from context to context, and may not even be well-defined. By building priors using the scheme described above, human observers are able to recover structure in a relatively reliable and flexible way, despite lacking crucial information. Like many heuristic interpretation schemes, this scheme is built on a concept of simplicity—here, on the idea that curvilinear processes tend to behave in a fairly constrained way, neither curving too much nor changing their curvature too quickly. What is most intriguing is the very precise way in which this simplicity principle translated into probabilistic terms, as a correlation. Progress in understanding other types of visual interpretation may be made if other heuristic inference principles can be fleshed out in a similarly concrete way.

REFERENCES

- Ashby, F. G. & Perrin, N. A. (1988). Toward a unified theory of similarity and recognition. *Psychological Review*, 95, 124–150.
- Atneave, F. (1954). Some informational aspects of visual perception. *Psychological Review*, 61, 183–193.
- Barlow, H. & Reeves, B. (1979). The versatility and absolute efficiency of detecting mirror symmetry in random dot displays. *Vision Research*, 19, 783–793.
- Binford, T. (1981). Inferring surfaces from images. *Artificial Intelligence*, 17, 205–244.
- Brookes, A. & Stevens, K. A. (1991). Symbolic grouping versus simple cell models. *Biological Cybernetics*, 65, 375–380.
- Caelli, T. M. & Julesz, B. (1978). On perceptual analyzers underlying visual texture discrimination: Part I. *Biological Cybernetics*, 28, 167–175.
- Dobbins, A., Zucker, S. W. & Cynader, M. S. (1988). Endstopped neurons in the visual cortex as a substrate for calculating curvature. *Nature*, 329, 438–441.
- Dobbins, A., Zucker, S. W. & Cynader, M. S. (1989). Endstopping and curvature. *Vision Research*, 29, 1371–1387.
- Dodwell, P. C. (1983). The lie transformation group model of visual perception. *Perception and Psychophysics*, 34, 1–16.
- Fahle, M. (1991). Parallel perception of vernier offsets, curvature, and chevrons in humans. *Vision Research*, 31, 2149–2184.
- Feldman, J. (1993). Perceptual models of small dot clusters. In *Proceedings of the DIMACS Workshop on Partitioning Data Sets*, 1993. Published as *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 19, 1995, 331–355.
- Feldman, J. (1996). Regularity vs. genericity in the perception of collinearity. *Perception*, 25, 335–342.
- Feldman, J. (1997a). Regularity-based perceptual grouping. *Computational Intelligence*, 13.
- Feldman, J. (1997b). The structure of perceptual categories. *Journal of Mathematical Psychology*, in press.
- 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.
- Foster, D. H. (1983). Visual discrimination, categorical identification, and categorical rating in brief displays of curved lines: implications for discrete encoding processes. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 785–807.
- Foster, D. H., Simmons, D. R. & Cook, M. J. (1993). The cue for contour-curvature discrimination. *Vision Research*, 33, 329–341.
- Glass, L. (1969). Moiré effects from random dots. *Nature*, 223, 578–580.
- Jepson, A. & Richards, W. A. (1992). What makes a good feature? In Harris, L. & Jenkin, M. (Eds), *Spatial vision in humans and robots*. Cambridge, U.K.: Cambridge University Press.
- Kendall, D. G. & Kendall, W. S. (1980). Alignments in two-dimensional random sets of points. *Advances in Applied Probability*, 12, 380–424.
- Koenderink, J. J. & Richards, W. A. (1988). Two-dimensional curvature operators. *Journal of the Optical Society of America A*, 5, 1136–1141.
- Kovacs, I. & Julesz, B. (1993). A closed curve is much more than an incomplete one: effect of closure in figure-ground segmentation. *Proceedings of the National Academy of Sciences USA*, 90, 7495–7497.
- Lowe, D. G. (1987). Three-dimensional object recognition from single two-dimensional images. *Artificial Intelligence*, 31, 355–395.
- Nosofsky, R. (1991). Typicality in logically defined categories: exemplar-similarity versus rule instantiation. *Memory and Cognition*, 19, 131–150.
- Parent, P. & Zucker, S. W. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11, 823–839.
- Prazdny, K. (1984). On the perception of glass patterns. *Perception*, 13, 469–478.
- Richards, W. A., Jepson, A. & Feldman, J. (1996). Priors, preferences, and categorical percepts. In Knill, D. & Richards, W. A. (Eds), *Perception as Bayesian inference*. Cambridge, U.K.: Cambridge University Press.
- Stevens, K. A. (1978). Computation of locally parallel structure. *Biological Cybernetics*, 29, 19–28.

- Takeichi, H. (1995). The effect of curvature on visual interpolation. *Perception*, 24, 1011–1020.
- Wagemans, J. (1993). Skewed symmetry: a nonaccidental property used to perceive visual forms. *Journal of Experimental Psychology: Human Perception and Performance*, 19, 364–380.
- Wagemans, J., Van Gool, L., Swinnen, V. & Van Horebeek, J. (1993). Higher-order structure in regularity detection. *Vision Research*, 33, 1067–1088.
- Watt, R. J. & Andrews, D. P. (1982). Contour curvature analysis: hyperacuity in the discrimination of detailed shape. *Vision Research*, 22, 449–460.
- Wilson, H. R. (1985). Discrimination of contour curvature: data and theory. *Journal of the Optical Society of America A*, 2, 1191–1199.
- Witkin, A. P. & Tenenbaum, J. M. (1983). On the role of structure in vision. In Beck, J., Hope, B. & Rosenfeld, A. (Eds), *Human and machine vision* (pp. 481–543). New York: Academic Press.
- Zucker, S. W. (1985). Early orientation selection: tangent fields and the dimensionality of their support. *Computer Vision, Graphics, and Image Processing*, 32, 74–103.
- Zucker, S. W. & Davis, S. (1988). Points and endpoints: a size/spacing constraint for dot grouping. *Perception*, 17, 229–247.

Acknowledgements—This research was supported by the Rutgers Center for Cognitive Science (RuCCS). I am grateful to Whitman Richards for many helpful discussions, to two anonymous reviewers for helpful comments on the manuscript, and to Janice Taylor for running the subjects.

APPENDIX

Derivation of the Axes of the Equiprobable Ellipse of a Bivariate Gaussian

We consider without loss of generality a bivariate gaussian with two

equal standard deviations of unity and correlation coefficient r . Each equiprobable contour has the form

$$\left(\frac{1}{1-r^2}\right)(a_1^2 + a_2^2 - 2ra_1a_2) = c^2 \quad (\text{A1})$$

for some constant probability c^2 . Now take the change of variables

$$\begin{aligned} u &= a_1 + a_2 \\ v &= a_1 - a_2, \end{aligned} \quad (\text{A2})$$

which rotates the coordinate frame by 45 deg. (A1) can be rewritten as

$$\left(\frac{1}{1-r^2}\right)\left[\frac{u^2 + v^2}{2} - \frac{r(u^2 - v^2)}{2}\right] = c^2, \quad (\text{A3})$$

which simplifies to

$$\left[\frac{1-r}{2(1-r^2)}\right]u^2 + \left[\frac{1+r}{2(1-r^2)}\right]v^2 = c^2. \quad (\text{A4})$$

The canonic form of an ellipse with principal axes α and β is

$$\frac{u^2}{\alpha^2} + \frac{v^2}{\beta^2} = c^2. \quad (\text{A5})$$

Hence, the equiprobable contour is an ellipse with principal axes

$$\alpha = \sqrt{\frac{2(1-r^2)}{1-r}} \quad (\text{A6})$$

$$\beta = \sqrt{\frac{2(1-r^2)}{1+r}}, \quad (\text{A7})$$

pointing in the a_1+a_2 and a_1-a_2 directions, respectively. The ratio of the two axes (the aspect ratio of the ellipse) is thus

$$\frac{\alpha}{\beta} = \sqrt{\frac{1+r}{1-r}}. \quad (\text{A8})$$