Contour extrapolation using probabilistic cue combination

Manish Singh
Psychology and Cognitive Science
Rutgers University—New Brunswick
Piscataway, NJ 08854
manish@ruccs.rutgers.edu

Jacqueline M. Fulvio
Psychology
New York University
New York, NY 10003
jmf384@nyu.edu

Abstract

A common approach to the problem of contour interpolation is based on the calculus of variations. The optimal interpolating contour is taken to be one that minimizes a given smoothness functional. Two important such functionals are total curvature (or bending energy) and variation in curvature. We analyzed contours extrapolated by human observers given arcs of Euler spirals that disappeared behind an occluding surface. Irrespective of whether the Euler spirals had increasing or decreasing curvature as they approached the occluding edge, visually-extrapolated contours were found to be characterized by decaying curvature. This curvature decay is modeled in terms of a Bayesian interaction between probabilistically-expressed constraints to minimize curvature and minimize variation in curvature. The analysis suggests that using fixed smoothness functionals is not appropriate for modeling human vision. Rather, the relative weights assigned to different probabilistic shape constraints may vary as a function of distance from the point(s) of occlusion. Implications are discussed for computational models of shape completion.

1. Introduction

1.1. Variational approaches to shape interpolation

A common approach to contour interpolation is based on the calculus of variations. The optimal interpolating contour is taken to be one that minimizes a given smoothness functional or energy term—which embodies a specific constraint on the shape of interpolating contours. Two prominent constraints stand out in computational work on contour interpolation. The first is that the interpolating contour must minimize the *total curvature* $\int \kappa^2 ds$ along its length. This constraint has its roots in the theory of elasticity, where the total curvature is referred to as a curve's "bending energy" [25]. Minimizing this energy leads to a class of curves known as *elastica*—curves that are as "straight" as possible given the boundary conditions imposed by the image contours and the requirement of smoothness [19, 27]. The second constraint involves the minimization of *variation in curvature* $\int \left(\frac{d\kappa}{ds}\right)^2 ds$. Rather than penalizing the presence of curvature, this constraint penalizes *changes* in curvature [3]. As a result, the contours tend towards being as close to circular arcs as the boundary conditions will allow, and generate a class of curves known as Euler spirals—that are characterized by a linear variation in curvature as a function of arc length [23]. There is indeed a history of work attributing a special status to circular arcs in contour interpolation and curve detection in noisy images. Ullman [36] modeled the shape of illusory contours as the combination of two circular arcs that respectively extrapolate the tangents of the two inducing contours, and intersect with continuous tangents. In the context of curve detection, Parent & Zucker [28] introduced the notion of edge co-circularity—i.e., tangency to a common circle—and used it to compute the strength of grouping between oriented image elements. The closer two edges are to being cocircular, the more strongly they are grouped.

1.2. Contour detection and integration

Although the contributions of the above two constraints in determining the *extended shapes* of contours interpolated by human vision have not been investigated, psychophysical work on contour detection and integration has provided evidence for the instantiation of *local* versions of these constraints in human contour perception. Observers' ability to visually integrate discrete local elements into extended contours has been found to deteriorate systematically with increasing curvature—defined in terms of the turning angles between successive local elements [11, 14, 16]. These results are generally taken to be consistent with an "association field" model in which the pattern of connection strengths between local orientation-tuned units is strongest when their preferred orientations are collinear, and decreases monotonically with the magnitude of the turning angle [14, 17]. This pattern of connection strengths is also

Proceedings of the 2006 Conference on Computer Vision and Pattern Recognition Workshop (CVPRW'06)
0-7695-2646-2/06 $20.00 © 2006 IEEE

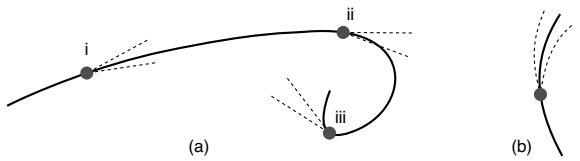

Figure 1. Generative models based on extrapolating tangent direction vs. curvature (adapted from August & Zucker [2]).

found to be consistent with the co-occurrence statistics of edge orientations along extended contours in natural images [8, 16]. Similarly, there is evidence for the local instantiation of minimization of *variation* in curvature in human contour integration. Observers' performance in contour integration tasks is best when the variance in the turning angles between successive local elements is minimal [11, 12, 29], namely, when the local elements are closest to being co-circular. Recent work on the statistics of natural images has also found that there is a prevalence of co-circular structure in natural images [8, 16, 31]. Moreover, a recent re-analysis of physiological data [5] by Ben-Shahar & Zucker [4] suggests that the association fields of individual orientation-tuned units in the primary visual cortex may in fact be tuned to different curvatures—with the "standard" shape of the association field being a description of the population average, rather than of each individual unit.

1.3. Generative models of contours

The above two constraints are naturally viewed as embodying two different generative models of contours—expressed as different distributions on "successive" oriented edges along contours. In discrete form, the minimization of curvature is consistent with a generative model in which the position of the "next" point (in angular terms, measured from the current contour direction) is characterized by a probability distribution centered on a turning angle of $\alpha = 0$ (i.e., "straight" is most likely) and falls off symmetrically with increasing magnitude of α [11, 13, 39]. The minimization of variation in curvature, on the other hand, is consistent with a generative model in which the position of the next point is characterized by an angular distribution centered on the previous turning angle (or a weighted average of the previous n turning angles)—so it tends to continue the estimated curvature of the contour.

The difference between the predictions of these two generative models is illustrated in the examples in Fig 1, adapted from August & Zucker [2]: Suppose an "ant" is walking along a curved path in thick fog. Based on where it has just come from, it makes predictions concerning where the path will go next. In utilizing the first generative model, it centers its prediction cone on the current tangent direction of the path. Fig 1a shows that this strategy works reasonably well—except when the path has high curvature, in which case the true path can easily lie outside the prediction cone. On the other hand, a strategy that takes into account the curvature of the contour as well, can do considerably better in capturing the true path, as illustrated in Fig 1b.

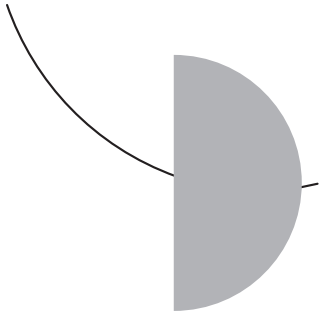

Figure 2. Stimulus configuration used to map out the visually-extrapolated shape of partly-occluded contours (adapted from Singh & Fulvio [33]).

A natural question is: which of these generative models better characterizes human vision? Clearly, further generative models may be considered as well—e.g., that take into account the rate of change of curvature of a contour, in addition to its tangent direction and curvature. On one hand, there is of course a benefit in utilizing generative models that incorporate higher-order derivatives, since these are likely to yield more accurate estimates of the true "path"—especially when it is highly complex. On the other hand, there is a clear cost associated with this strategy, since the computation of higher derivatives requires further computational resources and, perhaps more importantly, is increasingly more prone to noise. Determining how the human visual system resolves this tradeoff is largely an open question, and one that has important implications for computational models of human contour interpolation.

2. Visual Contour Extrapolation

In previous work [33], we measured the shape of visually-extrapolated contours for arcs of circles and parabolas that disappeared behind the vertical edge of a half-disk occluder (see Fig 2). Observers' task was to iteratively adjust the angular position of a short line probe around the circumference of the half disk, and its orientation, in order to optimize the percept of extrapolation of the partly-occluded contour. Using half disks with different radii allowed us to obtain measurements at multiple distances from the point of occlusion, and generate a relatively detailed representation of observers' extrapolated contours. Two main conclusions from that study were: (i) in addition to tangent direction, the visual system makes systematic use of curvature in extrapolating contour shape; observers' extrapolation curvature exhibited a linear depen-

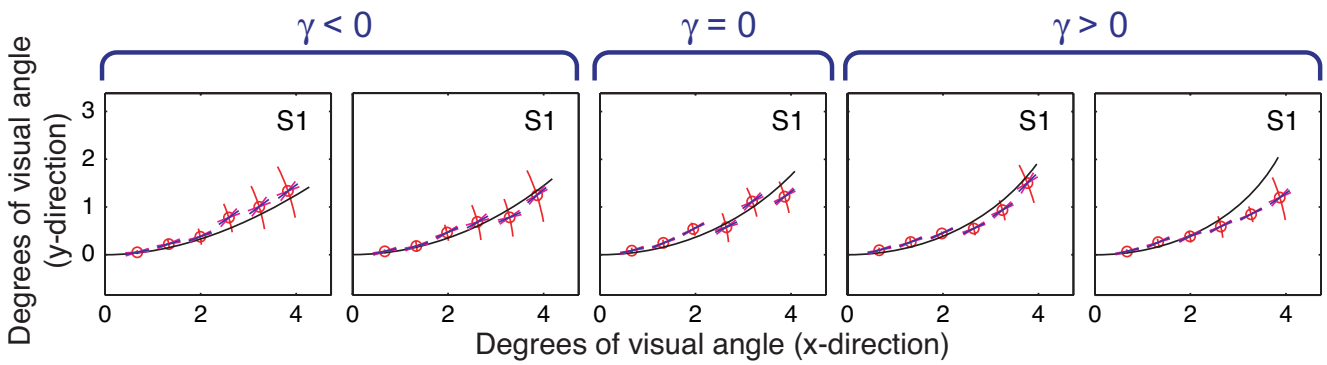

Figure 3. Extrapolation results for observer S1 for the 5 values of γ corresponding to the higher value of "initial" curvature κ_0 used. The complete data set comprises 40 such plots: 4 observers $\times$ 2 values of κ_0 $\times$ 5 values of γ.

dence on inducer curvature; and (ii) visually-extrapolated contours from arcs of circles and parabolas are characterized by decaying curvature with increasing distance from the point of occlusion.

2.1. A psychophysical study

In a recent psychophysical study [34], we sought to test the generality of these findings by using a wider and more flexible class of curves. The curves to be extrapolated were arcs of Euler spirals, which are characterized by a linear variation in curvature with arc length [23]. The use of these curves allowed us to manipulate the rate of change of curvature of the inducing contour independently from its curvature at the point of occlusion. In particular, the curvature of the contour could be either increasing or decreasing as it approached the occluding edge.

Two values of "initial" curvature κ_0 (i.e., at the point of occlusion) were crossed with 5 values of rate of change of curvature γ, thereby yielding 10 Euler-spiral contours. Two of the γ values were negative (curvature decreasing as the contour approaches the point of occlusion), one zero (arc of circle), and two positive (curvature increasing). Following the procedure used in our previous study [33], for each contour, observers made paired settings of angular position θ and orientation ϕ at each of 6 radial distances from the point of occlusion. On each trial, the contour to be extrapolated was presented at a random orientation, and either as concave up or concave down. Observers iteratively adjusted the angular position and orientation of the line probe in order to optimize the percept of smooth extrapolation.

4 observers performed adjustments in 8 experimental sessions each, preceded by one practice session. Each session consisted of 60 trials, with each trial requiring paired settings of angular position and orientation. This yielded, for each of the 10 contours used, 96 measurements per observer: 2 angular settings (angular position, orientation) $\times$ 6 radial distances $\times$ 8 repetitions [34]. This procedure yielded sufficient data per observer, so that the performance of various shape models could be compared in modeling each individual observer's data. All measurements were standardized with respect to the tangent direction of the inducing contour at the point of occlusion, i.e., as if the curve had been presented with its tangent horizontal at the point of occlusion, and with curvature concave up.

Fig 3 shows the extrapolation measurements for one of the observers, for 5 of the arcs of Euler spirals used (with the higher value of initial curvature, κ_0). At each radial distance, the mean setting of angular position is shown, with curved error bars depicting $\pm$ 1 standard deviation in θ, and the mean orientation is shown with error cones depicting $\pm$ 1 standard deviation in ϕ. The complete data set comprises 40 such plots: 4 observers $\times$ 10 Euler-spiral curves.

2.2. Shape modeling of extrapolation data

In order to model the shape of observers' extrapolated contours, we fitted an Euler-spiral model to the extrapolation measurements corresponding to each inducing contour (i.e., combined settings of angular position and orientation at the six different distances from the point of occlusion). To estimate the best-fitting Euler-spiral curve to an observer's extrapolation data, we computed the MLEs of the parameter pair (κ_0, γ)—i.e., the combination of "initial" curvature and rate of change of curvature that maximizes the likelihood function $\ell((\kappa_0, \gamma)|\mathcal{D})$, i.e., the probability $p(\mathcal{D}|(\kappa_0, \gamma))$ of obtaining a given data set $\mathcal{D}$ using an Euler-spiral curve with parameters κ_0 and γ. This likelihood model assumes that the observed settings of probe position and orientation result from the introduction of Gaussian noise to the ideal settings of angular position and orientation derived from the Euler-spiral generation process.[1]

[1]The appropriate model of noise for angular measurements is, strictly speaking, given by the von Mises distribution. However, given the small standard deviations in our data, the von Mises is very closely approximated by the Gaussian. (It converges to the Gaussian as $\sigma \to 0$ [26].)

Let $\mathcal{D} = \{\theta_r^i, \phi_r^i\}|_{i=1}^{8}|_{r=1}^{6}$ constitute an observer's extrapolation data set for a given Euler-spiral contour—consisting of settings of angular position and orientation, with eight repetitions ($1 \leq i \leq 8$) at each of the six radial distances ($1 \leq r \leq 6$). Let $\theta_e(\kappa_0, \gamma, r)$ and $\phi_e(\kappa_0, \gamma, r)$ be the expected values of angular position and orientation, respectively, based on an Euler-spiral generation process with "initial" curvature κ_0 and rate of change of curvature γ, and obtained at the r^{th} measurement distance from the point of occlusion. The likelihood of a parameter pair (κ_0, γ) under the set of extrapolation measurements $\{\theta_r^i, \phi_r^i\}|_{i=1}^{8}|_{r=1}^{6}$ is then given by:

$$\ell((\kappa_0,\gamma)|\{\theta_r^i,\phi_r^i\}) = \prod_{r=1}^{6}\prod_{i=1}^{8} BN\left(\begin{pmatrix} \theta_e(\kappa_0,\gamma,r) - \theta_r^i \\ \phi_e(\kappa_0,\gamma,r) - \phi_r^i \end{pmatrix}, \begin{pmatrix} \sigma_\theta^2(r) & \sigma_{\theta\phi}(r) \\ \sigma_{\theta\phi}(r) & \sigma_\phi^2(r) \end{pmatrix}\right) \quad (1)$$

where $BN((x,y)', Cov)$ is the bivariate normal distribution with mean $(0,0)'$ and covariance matrix Cov.

For each observer's data set corresponding to a given Euler-spiral contour, the parameters $(\hat{\kappa}_0, \hat{\gamma})$ that maximize the likelihood function were computed using unconstrained nonlinear optimization (Nelder-Mead simplex method). Fig 4 shows the mapping from Euler-spiral parameters (κ_0, γ) of the inducing contour, to the best-fitting parameters $(\hat{\kappa}_0, \hat{\gamma})$ of the observer's extrapolation measurements, depicted as a flow-field. The tails of the arrows correspond to the parameters of the inducing Euler-spiral contour, whereas the arrowheads correspond to the best-fitting extrapolation parameters of the Euler-spiral model.

A salient feature of these fits to the extrapolation data is that, in 39 of the 40 cases (4 observers $\times$ 10 Euler-spiral contours), the best-fitting value for the rate of change of curvature, $\hat{\gamma}$, is negative (see Fig 4). In order to test whether these estimates of extrapolation $\hat{\gamma}$ were reliably different from 0, the fits of the Euler-spiral model where compared with those of a circular-arc model (i.e., degenerate case of the Euler-spiral model with $\gamma = 0$), using the likelihood-ratio test for nested-model comparison. This test uses of the fact that, under the null hypothesis that the less full model is correct, the likelihood-ratio statistic,

$$\Delta = 2\log\left(\frac{L_e}{L_c}\right) \quad (2)$$

is distributed as a χ^2 with degrees of freedom equal to the difference in the number of parameters in the two models (here, 1), where L_e and L_c are respectively the maximized likelihood values under the two models. This test showed that 31 of the 39 negative estimates of extrapolation $\hat{\gamma}$ were significantly different from 0 at the .05 level, whereas the only positive estimate was not.

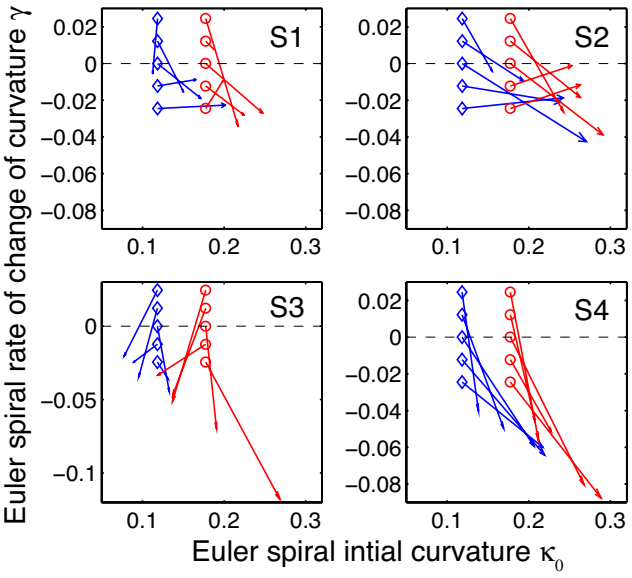

Figure 4. The mapping from parameters (κ_0, γ) of the inducing Euler-spiral contours, to the best-fitting parameters $(\hat{\kappa}_0, \hat{\gamma})$ of an Euler-spiral model to the extrapolation data. The tails of the arrows depict the inducing-contour parameters, and the arrows denote the visual-extrapolation parameters.

Thus visually-extrapolated contours are consistently characterized by monotonically-decreasing curvature, regardless of whether the curvature of the inducing contour is increasing or decreasing as it approaches the occluding edge. These findings thus extend our previous results on arcs of circles and parabolas [33], and indicate that curvature decay is a general property of visually-extrapolated contours—one that holds irrespectively of whether the inducing contour has increasing or decreasing curvature. They also suggest that the human visual system is unable to take into account the rate of change of curvature in extrapolating contour shape.

Fig 5 shows the dependence of estimated extrapolation rate of change of curvature on inducer curvature (at the point of occlusion), and inducer rate of change of curvature. In order to test this dependence formally, for each observer, the MLEs $\hat{\gamma}$ of extrapolation curvature were regressed on inducing Euler-spiral parameters, κ_0 and γ. The bilinear regression model revealed no statistically-reliable dependence of extrapolation rate of change of curvature on either inducer curvature or inducer rate of change of curvature, for any of the four observers. Observers are thus indeed unable to visually extrapolate the rate of change of curvature of a contour. When estimated extrapolation curvature $\hat{\kappa}_0$ was regressed on inducing Euler-spiral, three of the four observers exhibited a statistically significant dependence of extrapolation $\hat{\kappa}_0$ on inducer κ, indicating a systematic use of curvature in extrapolating contour shape.

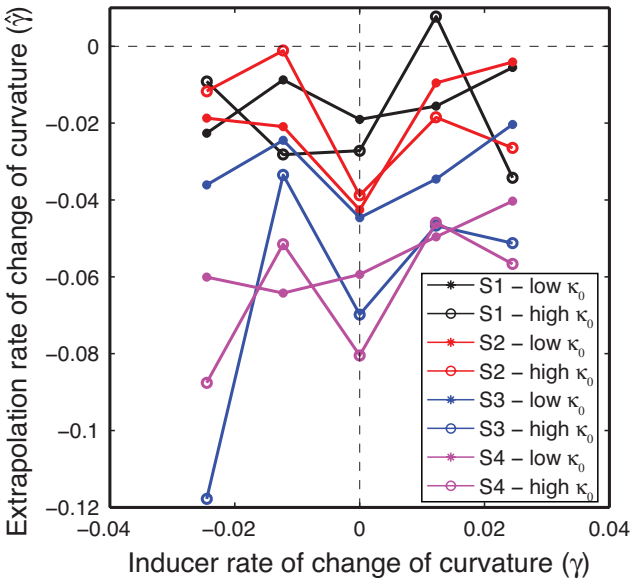

Figure 5. Dependence of extrapolation rate of curvature $\hat{\gamma}$ on Euler-spiral parameters κ_0 and γ.

2.3. Truncated Taylor series extrapolation

In our previous study, we found that a parabolic-arc model explained our extrapolation data better than a circular-arc model, irrespective of whether the inducing contours were arcs of circles or parabolas [33]. One interpretation of this result is that the visual system generates the extrapolated contour using an osculating parabola, based on the estimated curvature at the point of occlusion. (The special status of parabolas is also suggested by studies of interpolation of contours discretely sampled at dots [18, 37].) This essentially entails a strategy in which the visual system generates an extrapolating contour using a truncated Taylor series expansion based on the visible portion of the contour, with terms only up to the quadratic.[2]

In order to examine this possibility, we compared the fits of the Euler-spiral model to those of a parabolic model, in explaining the extrapolation data. Because the parabolic model is not nested within the Euler-spiral model, we used the Bayes Information Criterion [21] to compare the fits of the two models. The Bayes Information Criterion for a model is given by:

$$BIC = -2\log(L) + k\log(n) \tag{3}$$

where L is the maximized likelihood value under the model, k is the number of parameters in the model (1 for the parabolic model, 2 for the Euler-spiral model), and n is sample size ($= 96$, since each observer makes 48 settings of θ, and 48 settings of ϕ for each contour). In the BIC expression, the first term captures the goodness of fit to the data, whereas the second term constitutes a 'prior' penalty for more complex models, i.e., with more parameters. (The difference in the BIC values for two models asymptotically approaches twice the logarithm of the *Bayes factor*; see [21, 30]). The fits to the parabolic model were determined using a likelihood function with the same functional form as Eqn 1, but with the "expected" angular positions $\theta_p(\kappa_0, r)$ and orientations $\phi_p(\kappa_0, r)$ based on a parabola with curvature κ_0 at its vertex (the point of occlusion).

The BIC analysis showed that the Euler-spiral model was superior to the parabolic model in 31 of the 40 cases—thereby indicating that a parabolic model (or, equivalently, a Taylor series truncated beyond the quadratic term) does not provide the best characterization of contour shape extrapolated by human vision. In a majority of the cases, the increase in the goodness-of-fit to the data was sufficiently high to warrant the additional parameter in the Euler-spiral model. A comparison of the log-spiral model to the Euler-spiral model revealed the log-spiral model to be superior in 29 of the 40 cases. The superior fits of the log-spiral model are consistent with our previous findings [33], and indicate that a non-linear decay in curvature, asymptoting on zero, better characterizes the shape of visually-extrapolated contours than a linear decrease.

3. Bayesian interaction between minimizing curvature and variation in curvature

The above analyses indicate that, although human vision systematically takes into account the curvature of inducing contours, it does not use their rate of change of curvature in extrapolating their shape. Visually-extrapolated contours are consistently characterized by decaying curvature, regardless of whether the inducing contours have increasing or decreasing curvature. Moreover, the extrapolation rate of change of curvature exhibits no systematic dependence on the inducer rate of change of curvature.

We expand here on a Bayesian model outlined in [33], involving an interaction between the tendency to minimize curvature and the tendency to continue estimated curvature. The prior and the likelihood are both expressed as distributions on extrapolation curvature κ_{ext}. The prior captures the default expectation of the visual system—in the absence of any other information—that a contour will simply "go straight" [8, 11, 14, 16, 39]; and is expressed as a Gaussian distribution on extrapolation curvature, centered on 0:

$$p(\kappa_{ext}) \sim N(0, \sigma_{pr}) \tag{4}$$

for some σ_{pr}. This bias is clearly consistent with approaches that minimize total curvature along the length of the contour [19, 27]; it is, however, expressed as a probability distribution on local curvature. The likelihood bias captures the tendency toward co-circularity, namely, the

[2] We thank Steve Zucker for suggesting this hypothesis.

tendency to "continue" the curvature of the inducing contour estimated at the point of occlusion [8, 11, 16, 28, 29]. This constraint is expressed as a distribution on extrapolation curvature, centered on the estimated inducer curvature $\hat{\kappa}_i$; but is in spirit consistent with approaches that minimize variation in curvature along the length of the contour [23]. (One point of difference, however, is that the likelihood function takes into account the curvature of the inducing contour, whereas approaches that minimize variation in curvature typically do not. In particular, the total curvature variation is obtained by integrating over the interpolated portion of the contour only—i.e., not including the image contours. As a result, although the interpolated contour itself is smooth, curvature discontinuities are generally introduced at points where the interpolated contour meets the image contours; see [35].)

A key component of the model is the assumption that the continuation of inducer curvature is subject to systematically greater variability, with increasing distance from the point of occlusion. Specifically, a Weber-like dependence is assumed, such that the standard deviation increases linearly with distance d from the point of occlusion: $\sigma_{lik}(d) = \sigma^0_{lik} + md$, where σ^0_{lik} is the standard deviation when the gap size is zero (infinitesimally thin occluder). The likelihood function is thus given by:

$$\ell(\kappa_{ext}|\hat{\kappa}_i, d) \sim N(\hat{\kappa}_i, \sigma^0_{lik} + md). \tag{5}$$

The two constraints articulated above serve as probabilistic biases, or cues, to visual extrapolation. Their combination, via Bayes' Theorem, is given by:

$$p(\kappa_{ext}|\hat{\kappa}_i, d) = \frac{\ell(\kappa_{ext}|\hat{\kappa}_i, d) \cdot p(\kappa_{ext})}{p(\hat{\kappa}_i)} \tag{6}$$

Under the assumption that the prior and likelihood are Gaussian distributions, there exist standard formulas for the posterior [6]. In particular, the posterior is also a Gaussian with mean and variance given by:

$$\begin{aligned} \mu_{post}(\hat{\kappa}_i, d) &= \left(\frac{\mu_{pr}}{\sigma^2_{pr}} + \frac{\mu_{lik}}{\sigma^2_{lik}(\hat{\kappa}_i, d)}\right) \Bigg/ \left(\frac{1}{\sigma^2_{pr}} + \frac{1}{\sigma^2_{lik}(\hat{\kappa}_i, d)}\right) \\ \sigma^2_{post}(\hat{\kappa}_i, d) &= 1 \Bigg/ \left(\frac{1}{\sigma^2_{pr}} + \frac{1}{\sigma^2_{lik}(\hat{\kappa}_i, d)}\right) \end{aligned} \tag{7}$$

Interpreting these two biases or cues strictly as prior and likelihood is in fact not necessary. Treating the two distributions as probabilistic cues, or sources of information, the theory of cue combination (or sensor fusion) gives essentially the same expression for their combination. Specifically, the optimal combination of two stochastic "signals" (in the statistical sense of a minimum-variance unbiased estimator) is given by a weighted average of their expected values, with the weights being proportional to their respective *reliabilities* (i.e., reciprocals of their variances; see, e.g., [7, 24]). Thus, consistent with the expression for the expected extrapolation curvature above:

$$\hat{\kappa}_{ext}(d) = w_{pr} \cdot \mu_{pr} + w_{lik} \cdot \mu_{lik} \tag{8}$$

where $w_{pr} \propto 1/\sigma^2_{pr}$ and $w_{lik} \propto 1/\sigma^2_{lik}$.

Under the natural assumption that the continuation of estimated inducer curvature is subject to very little noise at the point of occlusion, we have $\sigma^0_{lik} \ll \sigma_{pr}$, and thus $w^0_{lik} \gg w_{pr}$. As a result, $\hat{\kappa}_{ext}(0) \approx \mu_{lik} = \hat{\kappa}_i$, i.e., the extrapolation curvature near the point of occlusion essentially equals the estimated inducer curvature. With increasing distance from the point of occlusion, the reliability of the 'likelihood' constraint to continue inducer curvature decreases systematically (because of an increase in its variance), whereas the variance of the prior remains constant. As a result, the curvature of the extrapolated contour is biased more and more toward the 'prior' curvature of zero; thus, extrapolation curvature decreases asymptotically to zero. The rate of curvature decay is modulated by the slope term m:

$$\hat{\kappa}_{ext}(d) = \hat{\kappa}_i \cdot \frac{\sigma^2_{pr}}{\sigma^2_{pr} + (\sigma^0_{lik} + md)^2}. \tag{9}$$

This probabilistic cue-combination model thus explains the decaying-curvature behavior of visually-extrapolated contours. Moreover, although the model makes systematic use of the curvature of the inducing contour, it involves no dependence on its rate of change of curvature. As our psychophysical results with Euler spirals show, this property is consistent with human visual extrapolation.

4. Implications for computational models of contour interpolation

The above experimental results have a number of implications for computational models of contour interpolation—at least insofar as they are meant to be consistent with biological vision. First and foremost is the role played by contour curvature in shape interpolation. Current models, whether they use geometric [22, 36], variational [19, 23, 27] or stochastic [38, 39] approaches, typically do not take into account the curvatures of the inducing contours—only their tangent directions at the points of occlusion. The results of our extrapolation studies show, however, that the human visual system makes systematic use of curvature in extrapolating contours. This indicates, as argued by Zucker and colleagues (e.g., [2, 4, 28]), that models of contour continuation must explicitly take into account contour curvature.

Second, variational approaches to shape completion generally begin with a particular smoothness functional, such as total curvature [19, 27] or variation in curvature [23], to be

minimized.[3] The results of our visual-extrapolation studies indicate, however, that the human visual system may not employ a fixed smoothness functional at all. Rather, the functional itself may vary—in particular, the relative weights assigned to various terms within the functional may vary—depending on distance from the point of occlusion. It is thus likely that the same is true of human contour *interpolation* as well since, after all, an interpolating contour must both (a) smoothly extrapolate the two inducing contours, and (b) smoothly connect the two extrapolants (see, e.g., [10, 36]). Thus, in the context of interpolation, the biases to continue the tangent directions and curvatures of the two inducing contours may be weighted differently at different locations—depending on the relative distances to the two points of occlusion (in addition to a general prior tendency to minimize curvature). The predictions of such a model have yet to be compared against detailed interpolation data from human observers, however.

Finally, the current results indicate that probabilistic interpolation models—wherein shape constraints are expressed as probability distributions over geometric variables—are likely to provide a richer and more appropriate class of models for human visual interpolation than those based on minimizing specific functionals. In particular, probabilistic models offer two main advantages. First, human interpolation data are intrinsically noisy, and probabilistic models allow the possibility of modeling not only the shape, but also the variability of human interpolation performance. Second, and more importantly, they allow for systematic changes in the shape of interpolating contours, by manipulating the variance of the relevant distributions. These variances are affected not only by factors such as distance from the point(s) of occlusion, but also by increased uncertainty in the stimulus itself. For instance, if an extremely short curved segment is presented, the curvature estimate is likely to be substantially less reliable. The effect that this increased variability has on predicted interpolation shape (via probabilistic cue-combination) can then be compared with its effect on contours interpolated by human vision.

This latter point is especially relevant in light of recent behavioral work that has demonstrated the influence of factors beyond local contour geometry (i.e., relative positions and orientations of inducing edges) on the shape of visually-interpolated contours. First, systematic shape differences have been shown to exist between corresponding partly-occluded (i.e., amodal) and illusory (i.e., modal) contours [1, 32]. Partly-occluded contours are perceived as being more angular than corresponding illusory contours (i.e., closer to the intersection of the two linear extrapolants)—a result that is consistent with the amodal completion have greater extrapolation strength than modal completion [32]. Second, in the context of stereoscopic illusory contours, the geometry of the *surface* they belong to has been shown to influence perceived illusory-contour shape. Illusory contours bounding locally convex surfaces are perceived to be smoother than those bounding locally concave surfaces [15] (see also [9] in the context of amodal completion); moreover, the influence of sign of curvature is modulated by the cross-axial width of the shape and its medial-axis geometry [15]. Probabilistic models permit sufficient flexibility to model such shape differences via the manipulation of the spread and/or central tendency of the relevant distributions, whereas it is more difficult to see how such influences could be accommodated within variational approaches.

[3] In some cases, a weighted mixture of the two is also used (e.g., [20]); however, it is nevertheless a fixed functional.

5. Conclusions

Our psychophysical experiments point to two basic findings on human visual extrapolation. First, visually-extrapolated contours are characterized by decaying curvature, irrespective of whether the contour to be extrapolated has increasing or decreasing curvature. Second, although the visual system makes systematic use of contour curvature in deriving the shape of extrapolated contours, it does not use rate of change of curvature. Both results are captured by a cue-combination model in which the tendency to minimize curvature and tendency to minimize variation in curvature are modeled as probability distributions on extrapolation curvature. In this model, the decaying-curvature behavior results from the systematic decrease in the reliability of the co-circularity (or "continue estimated curvature") cue with increasing distance from the point of occlusion. The model involves no dependence on a contour's rate of change of curvature.

These results have a number of implications for computational models of human contour interpolation: (i) interpolation models must utilize the curvature of inducing contours; (ii) the assumption of a *fixed* smoothness functional to be minimized does not appear to be valid for human vision; and (iii) probabilistic interpolation models provide a richer and more flexible class for modeling both the shape and the variability seen in human interpolation data. They may be especially useful in modeling the influences of factors beyond local contour geometry, such modal vs. amodal completion, and the role of the *surface* geometry.

References

[1] B. L. Anderson, M. Singh, and R. Fleming. The interpolation of object and surface structure. *Cognitive Psychology*, 44:148–190, 2002. 7

[2] J. August and S. W. Zucker. A Markov process using curvature for filtering curve images. In M. A. T. Figueredo, J. Zeurbia, and A. K. Jain, editors, *EMMCVPR 2001, Energy Minimization Methods in Computer Vision and Pattern*

Recognition, pages 497–512. Springer-Verlag, Sophia Antipolis, France, 2001. 2, 6

[3] H. G. Barrow and J. Tenenbaum. Interpreting line drawings as three-dimensional surfaces. *Artificial Intelligence*, 17:75–116, 1981. 1

[4] O. Ben-Shahar and S. W. Zucker. Geometrical computations explain projection patterns of long-range horizontal connections in visual cortex. *Neural Computation*, 16(3):445–476, 2004. 2, 6

[5] W. Bosking, Y. B. S. Zhang, and D. Fitzpatrick. Orientation selectivity and the arrangement of horizontal connections in the tree shrew striate cortex. *The Journal of Neuroscience*, 17(6):2112–2127, 1997. 2

[6] G. E. P. Box and G. C. Tiao. *Bayesian inference in statistical analysis*. John Wiley & Sons, New York, 1992. 6

[7] J. J. Clark and A. L. Yuille. *Data fusion for sensory information processing systems*. Kluwer, Boston, MA, 1990. 6

[8] J. H. Elder and R. M. Goldberg. Ecological statistics of Gestalt laws for the perceptual organization of contours. *Journal of Vision*, 2(4):324–353, 2002. 2, 5, 6

[9] C. Fantoni, M. Bertamini, and W. Gerbino. Contour curvature polarity and surface interpolation. *Vision Research*, 45:1047–1062, 2005. 7

[10] C. Fantoni and W. Gerbino. Contour interpolation by vector-field combination. *Journal of Vision*, 3(4):281–303, 2003. 7

[11] J. Feldman. Curvilinearity, covariance, and regularity in perceptual groups. *Vision Research*, 37(20):2835–2848, 1997. 1, 2, 5, 6

[12] J. Feldman. Bayesian contour integration. *Perception and Psychophysics*, 63(7):1171–1182, 2001. 2

[13] J. Feldman and M. Singh. Information along contours and object boundaries. *Psychological Review*, 112(1):243–252, 2005. 2

[14] D. Field, A. Hayes, and R. Hess. Contour integration by the human visual system: Evidence for a local "association field". *Vision Research*, 33(2):173–193, 1993. 1, 5

[15] J. M. Fulvio and M. Singh. Surface geometry influences the shape of illusory contours. *Acta Psychologica*, in press. 7

[16] W. S. Geisler, J. S. Perry, B. J. Super, and D. P. Gallogly. Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41(6):711–724, 2001. 1, 2, 5, 6

[17] S. Grossberg and E. Mingolla. Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review*, 92:173–211, 1985. 1

[18] A. Hon, L. Maloney, and M. Landy. The influence function for visual interpolation. *SPIE Conference Proceedings*, 3016:409–419, 1997. 5

[19] B. K. P. Horn. The curve of least energy. *ACM Transactions on Mathematical Software*, 9(4):441–460, 1983. 1, 5, 6

[20] M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *International Journal of Computer Vision*, 1:321–331, 1988. 7

[21] R. E. Kass and A. E. Raftery. Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795, 1995. 5

[22] P. J. Kellman and T. F. Shipley. A theory of visual interpolation in object perception. *Cognitive Psychology*, 23:141–221, 1991. 6

[23] B. B. Kimia, I. Frankel, and A. Popescu. Euler spiral for shape completion. *International Journal of Computer Vision*, 54(1/2):157–180, 2003. 1, 3, 6

[24] M. S. Landy, L. T. Maloney, E. Johnston, and M. Young. Measurement and modeling of depth cue combination: in defense of weak fusion. *Vision Research*, 35(3):389–412, 1995. 6

[25] A. E. H. Love. *Mathematical theory of elasticity*. University Press, Cambridge, 1927. 1

[26] K. V. Mardia. *Statistics of directional data*. Academic Press, London, 1972. 3

[27] D. Mumford. Elastica and computer vision. In C. L. Bajaj, editor, *Algebraic geometry and its applications*, pages 491–506. Springer-Verlag, New York, 1994. 1, 5, 6

[28] P. Parent and S. W. Zucker. Trace inference, curvature consistency and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11:823–839, 1989. 1, 6

[29] Z. Pizlo, M. Salach-Goyska, and A. Rosenfeld. Curve detection in a noisy image. *Vision Research*, 37(9):1217–1241, 1997. 2, 6

[30] G. Schwarz. Estimating the dimension of a model. *The Annals of Statistics*, 6:461–464, 1978. 5

[31] M. Sigman, G. A. Cecchi, C. D. Gilbert, and M. O. Magnasco. On a common circle: Natural scenes and Gestalt rules. *Proceedings of the National Academy of Sciences, USA*, 98(4):1935–1949, 2001. 2

[32] M. Singh. Modal and amodal completion generate different shapes. *Psychological Science*, 15:454–459, 2004. 7

[33] M. Singh and J. M. Fulvio. Visual extrapolation of contour geometry. *Proceedings of the National Academy of Sciences, USA*, 102(3):939–944, 2005. 2, 3, 4, 5

[34] M. Singh and J. M. Fulvio. Visual contour extrapolation: A Bayesian approach to good continuation. *Manuscript in preparation*, 2006. 3

[35] M. Singh and D. D. Hoffman. Completing visual contours: The relationship between relatability and minimizing inflections. *Perception and Psychophysics*, 61:636–660, 1999. 6

[36] S. Ullman. Filling-in the gaps: The shape of subjective contours and a model for their generation. *Biological Cybernetics*, 25:1–6, 1976. 1, 6, 7

[37] P. A. Warren, L. T. Maloney, and M. S. Landy. Interpolating sampled contours in 3D: analyses of variability and bias. *Vision Research*, 42(21):2431–2446, 2002. 5

[38] L. R. Williams and D. W. Jacobs. Stochastic completion fields: a neural model of illusory contour shape and salience. *Neural Computation*, 9:837–859, 1997. 6

[39] A. L. Yuille, F. Fang, P. Schrater, and D. Kersten. Human and ideal observers for detecting image curves. In S. Thrun, L. Saul, and B. Schölkopf, editors, *Advances in Neural Information Processing Systems, 16*, pages 59–70. MIT Press, Cambridge, MA, 2004. 2, 5, 6