Available online at www.sciencedirect.com
ScienceDirect
Vision Research 47 (2007) 783–798
Vision Research
www.elsevier.com/locate/visres

Bayesian contour extrapolation: Geometric determinants of good continuation

Manish Singh [a,*], Jacqueline M. Fulvio [b]

[a] Department of Psychology and Center for Cognitive Science, Rutgers University, New Brunswick Campus, 152 Frelinghuysen Road, Piscataway, NJ 08854-8020, USA
[b] Department of Psychology, New York University, New York, NY, USA

Received 12 May 2006; received in revised form 29 October 2006

Abstract

We investigated whether observers use rate of change of curvature in visually extrapolating contour shape. Arcs of Euler spirals with positive or negative rate of change of curvature γ (hence linearly increasing or decreasing curvature) disappeared behind the straight-edge of a half-disk occluder. Observers adjusted the position and the orientation of a line probe around the curved portion of the occluder to optimize the percept of extrapolation. These paired measurements were obtained at multiple distances from the point of occlusion in order to map out the extended shape of visually extrapolated contours. An Euler-spiral model was fit to the extrapolation data corresponding to each inducing contour. Maximum-likelihood estimates of extrapolation rate of change of curvature $\hat{\gamma}$ were consistently found to be negative, indicating that visually extrapolated contours are characterized by decaying curvature, irrespective of whether inducer curvature is increasing or decreasing as it approaches the occluder. Moreover, extrapolation $\hat{\gamma}$ was found to exhibit no systematic dependence on inducer γ. The results indicate that the visual system does not extrapolate rate of change of contour curvature. They support a Bayesian model of contour extrapolation, in which the decay in extrapolation curvature derives from an interaction between a likelihood bias to continue estimated contour curvature, and a prior bias to minimize contour curvature. Rate of change of curvature does not play a role.

Keywords: Contour; Shape; Extrapolation; Interpolation; Curvature; Occlusion; Bayes; Cue combination

1. Introduction

1.1. Contour completion

A fundamental problem faced by the visual system is the fragmentary nature of the retinal inputs. Large portions of object boundaries are often missing in the retinal images—either due to partial occlusion or because of insufficient local image contrast. Occlusion, in particular, poses a ubiquitous problem given the multiplicity of objects in the world and the loss of one spatial dimension during image projection (see Fig. 1a). In order to compute object and surface structure from fragmented image data, the visual system must solve two related problems. It must determine whether disparate image elements are in fact part of a single continuous contour/surface (the *grouping* problem); and, if so, what shape it has in the missing portions (the *shape* problem).

There has been a great deal of research on the *grouping problem* in a number of different contexts, including partly occluded contours (Fig. 1a), illusory contours (Fig. 1b), and discretely sampled contours—either sampled at dots (Fig. 1c) or at oriented line segments (Fig. 1d). Research on partly occluded contours and illusory contours has investigated the geometric conditions under which contour fragments belonging to distinct image regions are seen as belonging to a single perceptually completed contour

* Corresponding author. Fax: +1 732 445 6715.
E-mail address: manish@ruccs.rutgers.edu (M. Singh).

0042-6989/$ - see front matter © 2006 Elsevier Ltd. All rights reserved.
doi:10.1016/j.visres.2006.11.022

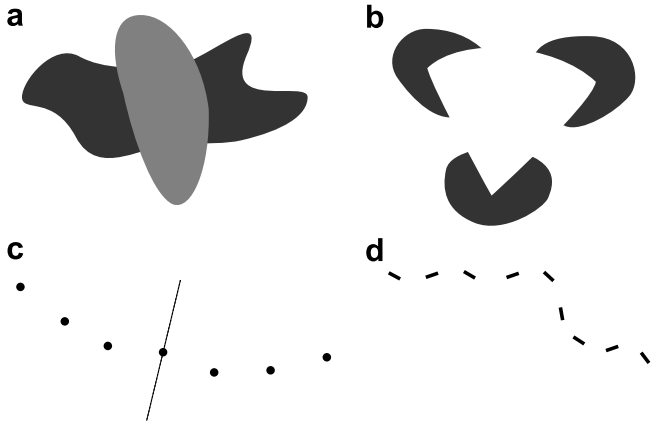

Fig. 1. Different contexts within which the problems of *grouping* and *shape* have been studied: (a) partial occlusion, (b) illusory contours, (c) contour integration from dot elements, and (d) contour integration from oriented elements. In each case, the visual system must determine (i) the likelihood that disparate local fragments are in fact part of continuous contour (the grouping problem), and (ii) what the shape of the contour is in the missing regions (the shape problem).

(Anderson, Singh, & Fleming, 2002; Fantoni & Gerbino, 2003; Guttman & Kellman, 2004; Heitger, von der Heydt, Peterhans, Rosenthaler, & Kübler, 1998; Kanizsa, 1979; Kellman & Shipley, 1991; Kubovy & Gepshtein, 2000; Ringach & Shapley, 1996; Singh & Hoffman, 1999; Takeichi, Nakazawa, Murakami, & Shimojo, 1995). Similarly, research on sampled contours (and contours embedded in noisy images) has investigated the geometric constraints that underly the grouping of discrete local elements—e.g., dots, line segments, Gabor patches—into extended contours (Caelli & Umansky, 1976; Feldman, 1997, 2001; Field, Hayes, & Hess, 1993; Geisler, Perry, Super, & Gallogly, 2001; Kovacs & Julesz, 1993; Pettet, McKee, & Grzywacz, 1998; Pizlo, Salach-Goyska, & Rosenfeld, 1997; Smits & Vos, 1987; Uttal, 1973), as well as how these constraints relate to the statistics of natural images (Elder & Goldberg, 2002; Geisler et al., 2001; Sigman, Cecchi, Gilbert, & Magnasco, 2001). The visual system's prior "expectations" about the distribution of relative orientations along contours have thus been investigated largely from the point of view of the grouping problem—i.e., based on whether, and how strongly, a set of discrete local elements group perceptually into the representation of a single extended contour.

By contrast, there has been relatively little psychophysical work on measuring the precise shapes of visually completed contours. The visual system faces a difficult problem in interpolating the shapes of "missing" portions of contours, because the problem is so highly underconstrained: Given a pair of inducing contours in an occlusion context, such as in Fig. 2a, there are infinitely many curves that could, in principle, smoothly interpolate between them. The fact that observers perceive only a very small subset of these possible solutions, entails that the visual system must impose strong geometric constraints in solving this problem.

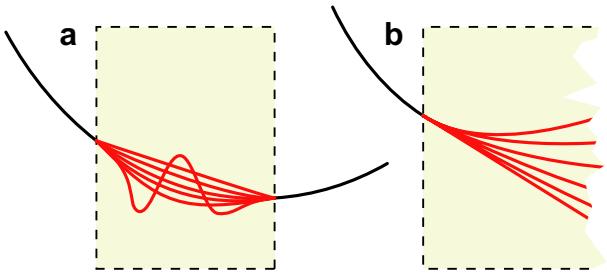

Fig. 2. Illustrating the highly unconstrained nature of the shape problem in the context of (a) interpolation between two contour segments, and (b) extrapolation of a contour segment.

Because visually completed contours are, by definition, synthesized by the visual system (being entirely absent in the images themselves), their shapes can reveal a great deal about the visual system's assumptions and constraints concerning contour shape—and hence about the underlying mechanisms of shape representation. The measurement of the detailed shape of partly occluded contours thus provides a unique, and largely untapped, opportunity to investigate the general shape constraints embodied in the visual processing of contours.

In the context of partly occluded contours, the contour shape interpolated between two contour segments (see Fig. 2a) has sometimes been measured using simple summary measures, such as the location of the extremal—e.g., highest—point along an interpolated contour segment (Guttman & Kellman, 2004; Takeichi, 1995), the estimated number of inflection points on an interpolated segment (Takeichi et al., 1995), ratings of its degree of 'roundedness' (Fantoni & Gerbino, 2003), or matching the overall degree of smoothness (Fulvio & Singh, 2006; Singh, 2004). Although these measures are sufficient to demonstrate that specific image variables influence the shape of visually interpolated contours, they do not provide a detailed representation of the perceived contour—which may then be analyzed for its *extended* shape. (See Anderson & Barth (1999), however, for a method that obtains *positional* measurements at multiple locations along a dynamic illusory contour.)

In the context of dot-sampled contours, a systematic investigation of interpolation performance has been carried out by Maloney, Landy, and colleagues (Hon, Maloney, & Landy, 1997; Warren, Maloney, & Landy, 2002, 2004). These researchers studied visual interpolation across gaps in linear and parabolic contours by having observers adjust the position of a probe dot so that it would appear to lie on the (invisible) smooth curve passing through the remaining dots (see Fig. 1c). They found that observers' settings were highly accurate, with no systematic bias away from the underlying linear and parabolic curves. Even in the case of parabolas rotated in 3D, the largest errors were only about 4 times larger than those expected based on a three-dot vernier-acuity task in the fronto-parallel plane (Klein & Levi, 1987; Warren et al., 2002). They found, however, an associated *cost of curvature* in that setting

variability was significantly greater for the parabola than for the line. Finally, two of their results made it clear that interpolation performance is determined relatively locally: (1) increasing the number of dots beyond 4 (to 6 or 8) did not significantly improve performance (Warren et al., 2002), and (2) interpolation settings of the probe dot were influenced only by the perturbation of the two nearest dots on each side; perturbing farther dots had hardly any influence on the settings (Hon et al., 1997; Warren, Maloney, & Landy, 2004). This locality of the *human visual spline* is consistent with Feldman's (1997) hypothesis that the visual system analyzes sampled contours through local windows containing four consecutive dots each.

1.2. Shape constraints

The current study focuses on the context of contour extrapolation (see Fig. 2a). Extrapolation was taken as a starting point for a number of reasons. First, extrapolation is a critical component of the more general problem of shape interpolation: an interpolating contour must both smoothly extrapolate each inducing contour, as well as smoothly connect the two extrapolants (Ullman, 1976; Fantoni & Gerbino, 2003). More importantly, extrapolation provides an ideal context within which to investigate the geometric properties that the visual system uses in continuing the shape of a contour. Specifically, our goal is to characterize the notion of "good continuation" of a contour in formal terms, by addressing the two following questions:

(1) What geometric properties (curvature, rate-of-change of curvature, etc.) of a contour does the visual system use in extrapolating its shape?
(2) How does it use and combine these variables to define the extended shape of a visually extrapolated contour?

Work in computational vision has proposed two shape constraints in solving the shape completion problem, that bear on these questions. The first is that an interpolating contour must minimize the *total curvature* $\int \kappa^2 \mathrm{d}s$ 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" (Euler, 1744/1952; Love, 1927). 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 physically specified edges, and the requirement of smoothness. Elastica have often been used in computer vision for interpolating contours between pairs of contour segments (Horn, 1983; Mumford, 1994). The second constraint involves the minimization of *variation in curvature* $\int (\frac{\mathrm{d}\kappa}{\mathrm{d}s})^2 \mathrm{d}s$. Rather than penalizing curvature per se, this constraint penalizes *changes* in curvature (e.g., Barrow et al., 1981; Kimia, Frankel, & Popescu, 2003; Singh & Hoffman, 1999). As a result, the contours tend locally toward being as close to circular as the boundary conditions will allow, and generate a class of curves known as Euler spirals—characterized by a linear variation in curvature as a function of arc length (Kimia et al., 2003). There is indeed a history of work attributing a special status to circular arcs in contour interpolation and curve detection in noisy images. Ullman (1976), for instance, modeled the shapes of illusory contours with the combination of two circular arcs that respectively extrapolate the tangents of the two inducing contours, and meet with continuous tangents. In the context of curve detection, Parent and Zucker (1989) 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.

The respective contributions of the above two constraints in determining the *extended shapes* of contours interpolated by the visual system have, to our knowledge, not been investigated. There is, however, psychophysical evidence for the instantiation of *local* versions of these constraints in the context of discretely sampled contours. In particular, observers' ability to visually integrate discrete local elements into contours is found to deteriorate systematically with increasing curvature—defined in terms of the turning angles between successive local elements (Feldman, 1997; Field et al., 1993; Geisler et al., 2001; Pettet et al., 1998; Uttal, 1973). These results are 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 increasing turning angle (Field et al., 1993; Grossberg & Mingolla, 1985). This pattern of connection strengths is also found to be consistent with the co-occurrence statistics of edge orientations along extended contours in natural images (Elder & Goldberg, 2002; Geisler et al., 2001). Similarly, there is evidence for the local instantiation of minimization of *variation* in curvature in human contour perception. Observers' performance in contour integration tasks is best when the variance in the turning angles between successive local elements is minimal (Feldman, 1997; Pizlo et al., 1997). Recent work on the statistics of natural images has also found that there is a prevalence of co-circular structure in natural images (Geisler et al., 2001; Sigman et al., 2001). Moreover, recent re-analysis of physiological data from Bosking, Zhang, and Fitzpatrick (1997) 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 (Ben-Shahar & Zucker, 2004).

The above two constraints are naturally viewed as embodying two different generative models of contours—expressed as different probability distributions on "successive" orientations along contours. In discrete form, the minimization of curvature is *locally* 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 (see, e.g., Feldman, 1997; Feldman & Singh, 2005; Yuille, Fang, Schrater, & Kersten, 2004). 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 a probability 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. 3, adapted from August and Zucker (2001): 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. 3a 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 path as well, can do considerably better, as illustrated in Fig. 3b. Further generative models that take into account higher-order derivatives may be naturally 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.

It is noteworthy that the generative models considered above make predictions only in the immediate neighborhood of a contour's current position. In the discrete form in which they have generally have been articulated (e.g., Feldman, 1997; Feldman & Singh, 2005; Yuille et al., 2004), a prediction is generally made for the position of the "next" point in the contour chain. In modeling contour-extrapolation performance by human vision, however, one would ideally like to characterize the visual system's "expectations" of where the contour will be at multiple distances from the current point—both in its immediate neighborhood, as well as further off.

In recent work, we measured the shape of extended portions of contours extrapolated from smooth contour segments—arcs of circles and parabolas—that disappeared behind a half-disk occluder (Singh & Fulvio, 2005). The

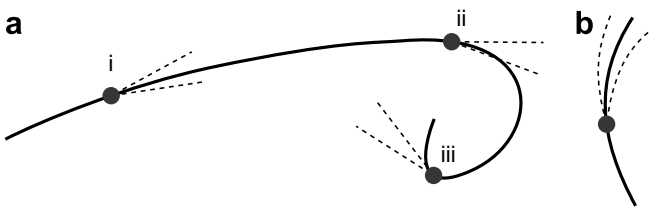

Fig. 3. Generative models of contours based on extrapolating (a) tangent direction only, and (b) tangent direction and curvature (adapted from August and Zucker, 2001).

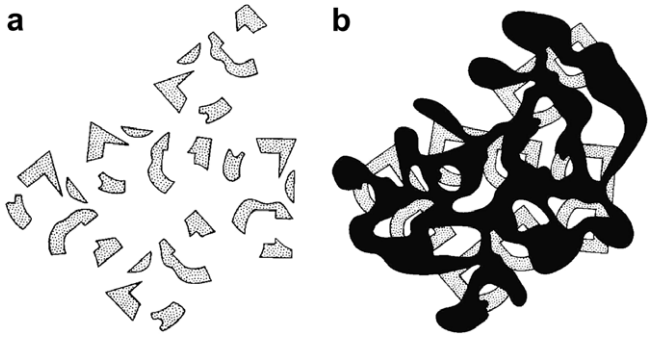

Fig. 4. Illustrating the role of occlusion in initiating mechanisms of visual completion (adapted from Bregman, 1981). The gray fragments are identical in (a) and (b), but are more easily completed and recognized in (b).

use of smooth (rather than discretely sampled) contours, and the use of occlusion (rather than, say, a contour simply coming to an end) both served to trigger mechanisms of visual completion (see, e.g., Nakayama, He, & Shimojo, 1995)—and hence generate a vivid percept of extrapolated-contour shape. (A striking demonstration, due to Bregman (1981), of the efficacy of occlusion cues in triggering mechanisms of visual completion is shown in Fig. 4.) Observers iteratively adjusted 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 inducing contour (see Fig. 5). Measurements were obtained at multiple distances from the point of occlusion by manipulating the radius of the half-disk occluder. This procedure generated a relatively detailed representation of an observer's extrapolated contour. The use of the half disks ensured that distance of the probe from the point of occlusion was preserved, as observers adjusted its position.[1]

For the purposes of the current paper, the most relevant results from that study were:

(1) observers made systematic use of inducer curvature in extrapolating their shape; for each observer, extrapolation curvature increased linearly with inducer curvature; and
(2) contours extrapolated from arcs of circles and parabolas were characterized by decaying curvature, with increasing distance from the point of occlusion.

The result that the visual system uses curvature in extrapolating contour shape is in itself significant in light of the fact that most current models of contour interpolation, in both human and computer vision, use only inducer orientation at the point of occlusion, but not inducer curvature. This is especially surprising given that the need to

[1] Moreover, as our data with linear inducers showed, the use of half disks as occluders also had the benefit of essentially eliminating the Pöggendorf illusion in our displays (with inducer orientations between 15 and 45 degrees from the horizontal).

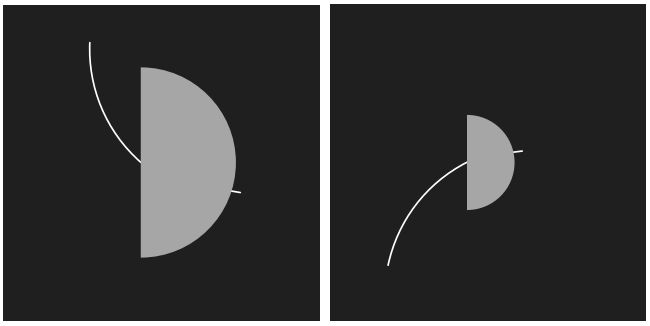

Fig. 5. The basic stimulus configuration used to measure the shapes of visually extrapolated contours. A curved inducing contour disappears behind the straight edge of half-disk occluder. Observers iteratively adjust the angular position and the orientation of a line probe protruding from behind the opposite, curved, side of the occluder, in order to optimize the percept of extrapolation. Obtaining measurements with half disks of multiple radii allows one to build up a detailed representation of the extended shape of a visually extrapolated contour.

incorporate curvature into these models has long been argued for (August & Zucker, 2001; Ben-Shahar & Zucker, 2004; Dobbins, Zucker, & Cynader, 1987; Parent & Zucker, 1989).

Given our goal of characterizing "good continuation" of a smooth contour in terms of (i) the geometric variables that are used in visually extrapolating contour shape, and (ii) characterizing the way in which these variables are used and combined to define its extrapolated shape, a natural question that arises is: What role is played by the rate of change of curvature? Extrapolation performance would of course be enhanced if the visual system took into account the rate of change of curvature as well. (Consider, for instance, a model in the context of August & Zucker's ant walking along a curved path, that bases its prediction on orientation, curvature, *and* rate of change of curvature; see Fig. 3.) On the other hand, there is a clear cost associated with this strategy as well, since the computation of higher derivatives requires further computational resources, and is increasingly more prone to noise. Determining how the human visual system resolves this tradeoff is largely an open question, and one with important implications for computational models of human shape completion.

In order to address this question, the current study measured visual extrapolation performance using arcs of Euler spirals—characterized by a linear variation of curvature with arc length. By manipulating the rate of change of curvature of the Euler spirals, to include both positive and negative rates (i.e., linearly increasing and decreasing curvatures), we sought to understand the role played by rate of change of curvature in the visual extrapolation of smooth contours.

2. Experiment

The basic stimulus configuration used is shown in Fig. 5. A curved inducing contour disappears behind the straight edge of a half-disk occluder. An oriented line probe protrudes from behind the opposite, curved, portion of the half disk. Observers perform paired adjustments—iteratively adjusting the position of the probe along the circumference of the half disk, and its orientation—in order to optimize the percept of continuation of the inducing contour. Measurements are obtained with half disks of multiple radii in order to generate a detailed representation of the visually extrapolated shape for each inducing contour.

2.1. Methods

2.1.1. Observers

Four observers at Rutgers University, with normal or corrected-to-normal visual acuity, participated in the experiment: three naïve observers and author JF (observer O2).

2.1.2. Stimuli and design

Stimulus displays comprised three components: an inducing contour, a half-disk occluder, and an adjustable line probe (see Fig. 5). These were presented against a homogeneous black background.

The inducing contours were arcs of Euler spirals, characterized by a linear variation of curvature with arc length s:

$$\kappa(s) = \kappa_0 + \gamma s \tag{1}$$

where κ_0 is the curvature at the point of occlusion, and γ is the rate of change of curvature in the direction approaching the occluder. (A positive value of γ thus means that the curvature of the inducing contour is linearly increasing as it approaches the vertical occluding edge of the half disk.) 10 inducing contours were used, generated by crossing 2 values of κ_0 with 5 values of γ. The two values of κ_0 used were: 0.118 and 0.178 deg^{-1}. The five values of γ used were: -0.0245, -0.01225, 0, 0.1225 and 0.0245 deg^{-2}. Thus, two of the γ values correspond to decreasing curvature of the inducing contour, one to constant curvature (i.e., an arc of a circle), and two to increasing curvature. The arc lengths of the visible inducing contours were fixed at 4.56 degrees of visual angle. The inducing contour segments were presented at random orientations—assigned by first sampling a magnitude from a uniform distribution between 15° and 45°, and then giving it a random sign. (The orientation of an inducing contour is specified by its tangent direction at the point of occlusion.) Half of the experimental sessions presented the inducing contours as concave up, the other half presented them as concave down.

The occluder was a mid-gray half disk (luminance = 27 cd/m^2), with its straight-edge vertical. The mid-point of this straight edge served as the point of occlusion, i.e., where the inducing contour disappeared behind the half disk. This ensured that the distance of the line probe from the point of occlusion was preserved as its posi-

tion was adjusted around the curved portion of the half disk's circumference. The radius of the half disk could take one of six values: 0.68, 1.35, 2.03, 2.7, 3.38, or 4.06 degrees of visual angle.

The oriented probe was a line segment, whose midpoint—serving as its pivot point—was constrained to lie along the curved portion of the half disk's circumference. The probe was placed behind the occluder, so that only the portion of the probe not overlapping with the half disk was visible (10.2 min of arc). The probe thus appeared to protrude from underneath the curved portion of the half disk. Both the angular position and the orientation of the line probe were to be adjusted by the observer.

The inducing contour segment and the line probe were white (luminance = 78 cd/m^2) and two-pixels thick (1.9 min of arc). They were anti-aliased to produce a smooth appearance, at the resolution of $\frac{1}{4}$ of a pixel.

2.1.3. Procedure

Each observer performed adjustments in 9 sessions, with the first session being practice. Within a session, each inducing contour was presented once with each half-disk radius. As noted above, there were 10 different inducing contours (2 values of $\kappa_0 \times 5$ values of γ) and 6 half-disk radii. Each session thus contained 60 trials, with each trial requiring paired adjustments of angular position and orientation of the line probe. Four of the experimental sessions for each observer presented the inducing contours as concave up, the other four presented them as concave down. Their order was counterbalanced.

Each trial began with one of the inducing contours presented at a randomly assigned orientation at the point of occlusion (as described above), and the line probe positioned either near the top or the bottom of the occluding half disk (polar angle $\theta^* = \pm 80°$) with horizontal orientation ($\phi^* = 0°$). Using a trackball, observers adjusted the position of the line probe around the circumference of the half disk. Pressing the space bar allowed them to toggle to adjusting the orientation of the line probe, while maintaining its positional setting. Observers iteratively adjusted the angular position and the orientation of the line probe in this manner, in order to optimize the percept of extrapolation. They pressed the trackball button when done with a given inducing contour. A question appeared at the bottom of the screen, asking them to verify that they were ready to move on to the next trial. Responding in the negative allowed them to return to their setting, and continue adjustment.

The experiment was programmed in MATLAB using the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997). The stimulus displays were presented on a high-resolution 22-in. monitor (*Lacie Blue*), under conditions of low ambient illumination. Observers viewed the stimuli from a distance of 102.5 cm, their viewing position fixed by means of a head-and-chin rest.

2.2. Results

Each observer's raw data comprised 8 paired settings of angular position θ^* and orientation ϕ^* for each of the 10×6 combinations of inducing contour and half-disk radius. These measurements were first standardized by transforming them into a single, canonical coordinate frame—one that treats the inducing contour as if it were presented horizontally at the point of occlusion, and with curvature concave up (see Fig. 6). The standardized settings thus measure the polar angle θ and orientation ϕ of the adjusted line probe relative to the tangent direction of the inducing contour at the point of occlusion—with positive values of θ and ϕ corresponding to the same sign of curvature as the inducing contour. (In the standardized coordinates, $\theta = \phi = 0$ corresponds to a linear extrapolation of the tangent direction of the inducing contour at the point of occlusion.) The measurements were collapsed over sessions presenting the inducing contours as concave-up and concave-down, as preliminary analysis revealed no systematic difference between them.

The extrapolation results for each observer are shown in Fig. 7. The standardized settings of angular position θ and orientation ϕ of the line probe, at each of the six radial distances, are shown embedded in the Cartesian plane with the point of occlusion at the origin. The mean angular-position settings are depicted by the red circles—with the error bars denoting standard deviations at each radial distance. The mean orientation settings are depicted by the blue oriented line segments passing through the mean position settings—with the error cones depicting ± 1 standard deviation around each mean orientation setting. Also shown on the plots are the "true" extrapolants—i.e., the extensions of the arcs of Euler spirals used to generate the inducing contours.

In order to model the shape of observers' extrapolated contours, we fit an Euler-spiral model to the extrapolation data corresponding to each inducing contour (i.e., combined settings of position and orientation at the six different distances from the point of occlusion). To

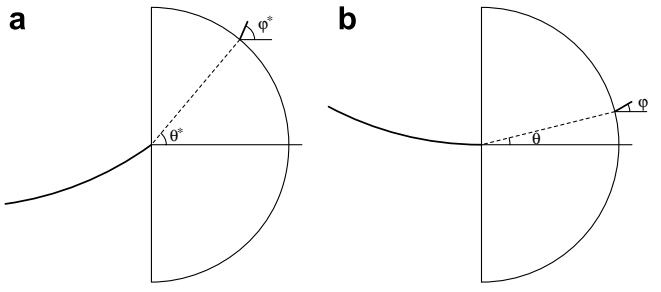

Fig. 6. (a) Raw measurements of the angular position θ^* and orientation ϕ^* of the line probe. (b) Standardized measurements of angular position θ and orientation ϕ, transformed into a coordinate frame in which the inducing contour is horizontal at the point of occlusion, with curvature concave up.

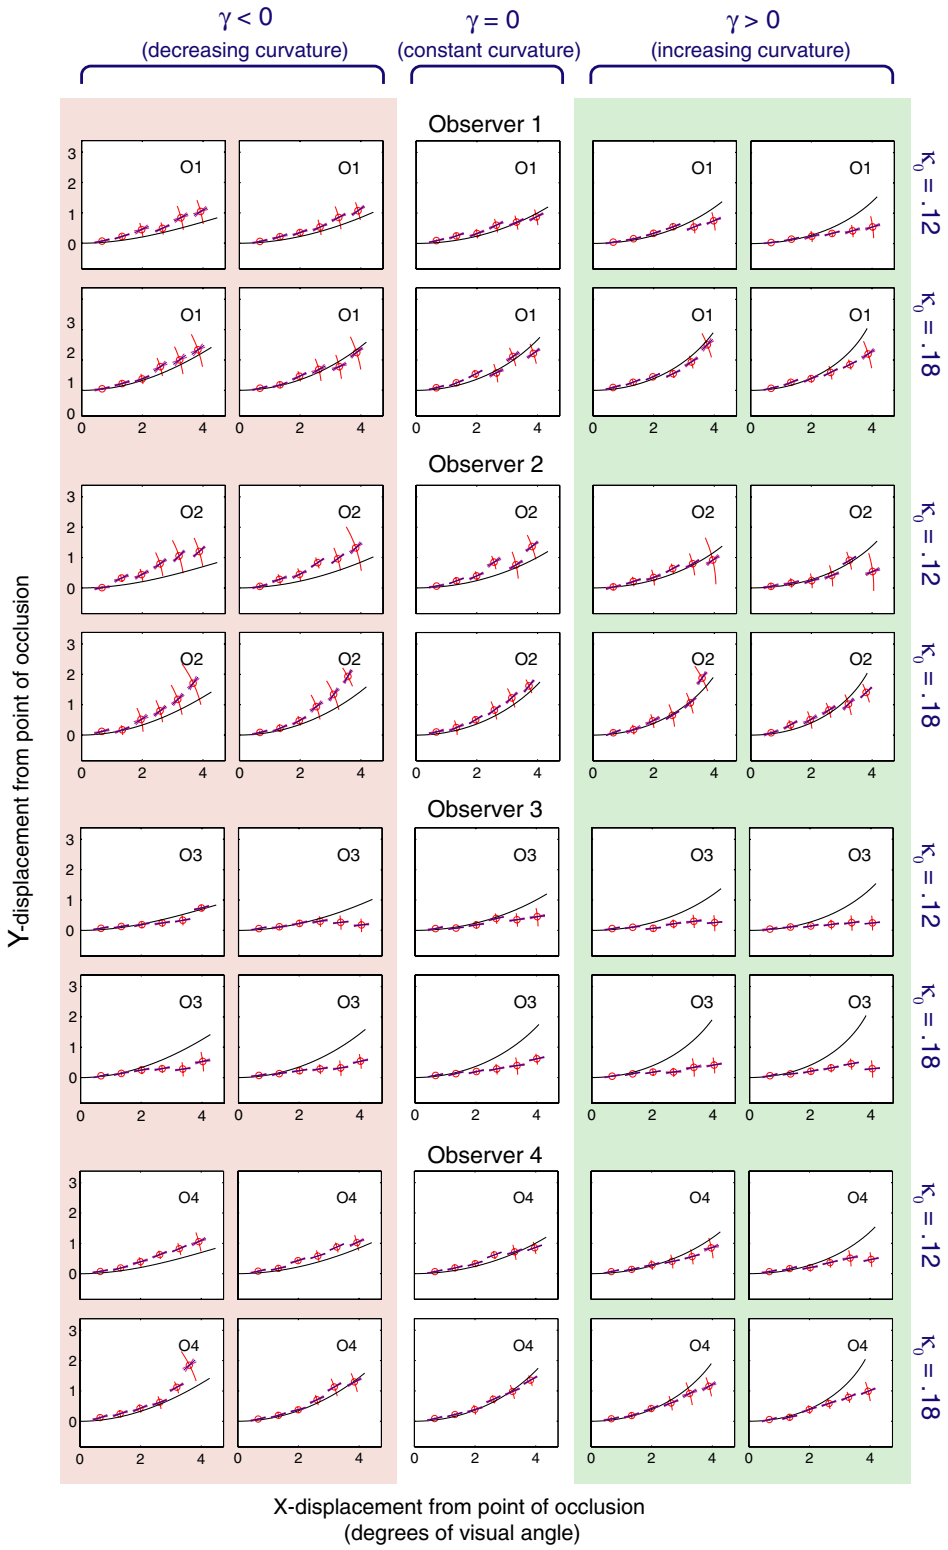

Fig. 7. Extrapolation results for the four observers. The top row for each observer corresponds to the lower initial curvature, the bottom row to the higher initial curvature. The two leftmost columns correspond to negative inducer γ (decreasing curvature as the inducing contour approaches the occluding surface), the two rightmost columns to positive inducer γ, and the middle column to circular arcs.

estimate the best-fitting Euler-spiral curve to an observer's extrapolation data, we computed the maximum-likelihood estimates (MLEs) of the parameter pair (κ_0, γ)—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 model with parameters κ_0 and γ.[2] 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.[3]

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 ($1 \leqslant i \leqslant 8$) for each of the six radial distances ($1 \leqslant r \leqslant 6$). 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 extrapolation generation process with "initial" curvature κ_0 and rate of change of curvature γ, and obtained at the rth 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:

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

where $BN((x\ y)', \mathrm{Cov})$ 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) as implemented in Matlab's fminsearch function. Fig. 8 shows the mapping from the parameters (κ_0,γ) of the inducing contour, to the best-fitting parameters $(\hat{\kappa}_0, \hat{\gamma})$ to 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 × 10 Euler-spiral contours), the best-fitting value for the rate of change of curvature, $\hat{\gamma}$, is negative (see Fig. 8)—indicating that the visually extrapolated contours consistently have decreasing curvature with increasing distance from the point of occlusion. 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 statistic:

[2] The initial orientation of the Euler-spiral extrapolant was constrained to match the orientation of the inducing contour at the point of occlusion.

[3] 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 (SDs <10° for θ, SDs <15° for ϕ), the von Mises distribution is very closely approximated by the Gaussian. (Indeed, the von Mises converges to the Gaussian as $\sigma \to 0$; see, e.g., Mardia & Jupp (2000).)

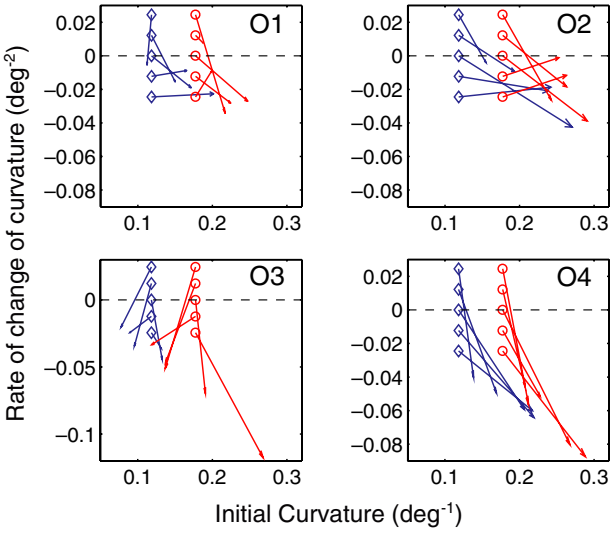

Fig. 8. The mapping from parameters (κ_0, γ) of the inducing Euler-spiral contours, to the best-fitting parameters $(\hat{\kappa}_0, \hat{\gamma})$ of an Euler-spiral model to the extrapolation data. The tails of the arrows depict the inducing-contour parameters, and the arrows denote the visual-extrapolation parameters.

$$\Delta = 2\log\left(\frac{p(\mathcal{D}|\hat{\kappa},\hat{\gamma},\mathcal{M}_{es})}{p(\mathcal{D}|\hat{\kappa},\mathcal{M}_c)}\right) \tag{3}$$

where $p(\mathcal{D}|\hat{\kappa},\hat{\gamma},\mathcal{M}_{es})$ and $p(\mathcal{D}|\hat{\kappa},\mathcal{M}_c)$ are the maximized likelihood values under the Euler-spiral model and circular-arc model, respectively—i.e., the respective likelihoods attained at the MLE parameter values under each model. Because $\mathcal{M}_c$ is nested within $\mathcal{M}_{es}$, the ratio of the likelihoods cannot be smaller than 1, and therefore the statistic Δ cannot be negative. In order to set a cut-off value, we use the fact that, under the null hypothesis that model $\mathcal{M}_c$ is correct, Δ is asymptotically distributed as a χ^2 with 1 degree of freedom (= the difference in the number of parameters between the two models; see Mood, Graybill, & Boes, 1974). This test showed that 31 of the 39 negative estimates of extrapolation $\hat{\gamma}$ were significantly different from 0 at the .05 level (the only positive estimate was not); see Table 1.

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. This extends our previous findings (Singh & Fulvio, 2005), showing 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. It also suggests that the human visual system does not take into account the rate of change of curvature in extrapolating the shape of contours.

Fig. 9 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, each observer's MLEs of extrapolation curvature, $\hat{\gamma}$, were regressed on inducing Euler-spiral parameters κ_0 and γ. The bilinear regression model revealed no statistically reliable

Table 1
Values of the log likelihood-ratio statistic for the Euler-spiral model $\mathcal{M}_{es}$ against the circular-arc model $\mathcal{M}_c$ (i.e., degenerate case of the Euler-spiral model with $\gamma = 0$)

γ	$\kappa_0 = 0.118$				$\kappa_0 = 0.178$			
	O1	O2	O3	O4	O1	O2	O3	O4
−0.0245	9.63*	2.41	28.17*	64.80*	1.53	1.70	366.92*	140.24*
−0.0123	1.53	5.24*	19.99*	81.35*	8.29*	2.20	26.87*	90.23*
0	7.35*	51.45*	56.96*	104.32*	20.09*	22.72*	190.75*	173.28*
+0.0123	9.37*	1.82	34.73*	77.45*	3.40	13.02*	44.47*	36.81*
+0.0245	2.30	0.27	18.47*	31.94*	27.19*	6.24*	86.64*	75.57*

Cases in which the null hypothesis of the circular model is rejected at the .05 level—and thus the $\hat{\gamma}$ estimate is significantly different from 0—are marked with asterisks.

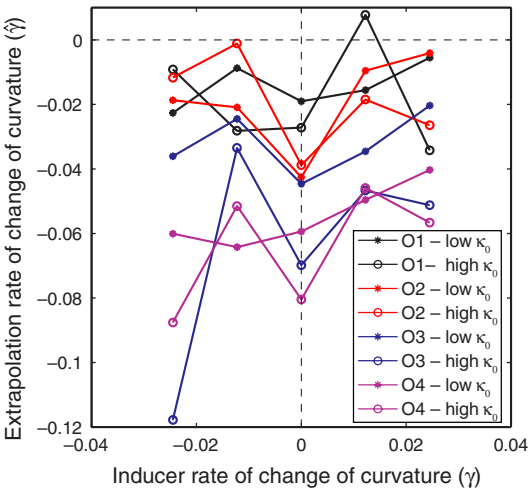

Fig. 9. Dependence of extrapolation rate of curvature $\hat{\gamma}$ on Euler-spiral parameters κ_0 and γ.

dependence of extrapolation rate of change of curvature on either inducer curvature or inducer rate of change of curvature, at the .05 level, for any of the four observers (see Table 2). The lack of a statistically reliable dependence of extrapolation $\hat{\gamma}$ on inducer γ, in particular, indicates that observers are unable to extrapolate the rate of change of curvature of a contour.

When estimated extrapolation curvature $\hat{\kappa}_0$ was regressed on inducing Euler-spiral parameters, three of the four observers exhibited a statistically reliable dependence of extrapolation $\hat{\kappa}_0$ on inducer κ_0 ($p < .05$ for three observers, $p = 0.066$ for the fourth (O2); see Table 3), consistent with our previous result that the visual makes systematic use of curvature in extrapolating a contour's shape (Singh & Fulvio, 2005). There was also a slight trend for the extrapolation $\hat{\kappa}_0$ to decrease with increasing inducer γ, but it was statistically reliable for only one observer.

Table 2
Regression parameters for estimated extrapolation rate of change of curvature $\hat{\gamma}$ on inducer curvature κ at the point of occlusion, and rate of change of curvature γ

Regression parameters	$\hat{\gamma}$			
	O1	O2	O3	O4
κ_{slope}	−0.066	−0.003	−0.537	−0.1638
95% CI	(−0.42, 0.29)	(−0.40, 0.39)	(−1.12, 0.04)	(−0.47,0.14)
γ_{slope}	0.053	−0.026	0.576	0.496
95% CI	(−0.55, 0.66)	(−0.70, 0.64)	(−0.42, 1.57)	(−0.02,1.01)
R^2	0.033	0.0012	0.488	0.493
p	0.890	0.996	0.096	0.093

Table 3
Regression parameters for estimated curvature $\hat{\kappa}$ at the point of occlusion on inducing contour parameters κ and γ

Regression parameters	$\hat{\kappa}$			
	O1	O2	O3	O4
κ_{slope}	0.927	0.719	1.098	0.854
95% CI	(0.28, 1.57)	(−0.05, 1.49)	(0.07, 2.12)	(0.35, 1.36)
γ_{slope}	−0.796	−1.017	−1.412	−1.557
95% CI	(−1.90, 0.31)	(−2.33, 0.30)	(−3.16, 0.34)	(−2.42, −0.69)
R^2	0.673	0.539	0.590	0.829
p	0.020	0.066	0.044	0.002

2.3. Truncated Taylor series extrapolation

In our previous study, we found that a parabolic-arc model explained the extrapolation data better than a circular-arc model, irrespective of whether the inducing contours were arcs of parabolas or circles (Singh & Fulvio, 2005). 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. This would entail a strategy whereby it 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. Such a strategy would be consistent with our current finding that the visual system uses the first and second-order derivatives of a contour—orientation and curvature, assuming arc-length parametrization—but not higher-order derivatives.

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 Bayesian model selection to compare the fits of the two models. Specifically, we used the Schwarz criterion, which

Table 4
Schwarz-criterion values for the Euler-spiral model $\mathcal{M}_{es}$ against the parabolic model $\mathcal{M}_{p}$

γ	$\kappa_0 = 0.118$				$\kappa_0 = 0.178$			
	O1	O2	O3	O4	O1	O2	O3	O4
−0.0245	−1.51	3.82*	9.78*	20.91*	0.94*	18.16*	155.55*	26.82*
−0.0123	−0.79	2.14*	7.29*	28.65*	−2.03	16.52*	9.94*	15.81*
0	−1.55	−2.28	23.80*	39.75*	1.35*	1.36*	82.06*	50.06*
+0.0123	−1.22	1.13*	14.78*	32.03*	15.29*	31.61*	17.85*	5.68*
+0.0245	−1.47	−1.81	6.70*	12.18*	0.94*	−0.62*	39.87*	25.48*

Cases in which the increase in the goodness of fit of the Euler-spiral model is sufficiently high to offset its additional parameter, are indicated by positive values (marked with asterisks).

provides an asymptotic approximation to the logarithm of the Bayes factor, i.e., the ratio of marginal likelihoods $p(\mathcal{D}|\mathcal{M}_{es})/p(\mathcal{D}|\mathcal{M}_{p})$ under the two models (see Kass & Raftery, 1995; Schwarz, 1978). The Schwarz criterion is given by:

$$S = \log\left(\frac{p(\mathcal{D}|\hat{\kappa}_0, \hat{\gamma}, \mathcal{M}_{es})}{p(\mathcal{D}|\hat{\kappa}_0, \mathcal{M}_{p})}\right) - \frac{1}{2}(k_{es} - k_{p})\log(N) \quad (4)$$

where $p(\mathcal{D}|\hat{\kappa}_0, \hat{\gamma}, \mathcal{M}_{es})$ and $p(\mathcal{D}|\hat{\kappa}_0, \mathcal{M}_{p})$ are respectively the maximized likelihood values under the two models (the likelihoods corresponding to the MLE parameter values), k_{es} and k_{p} are their respective numbers of parameters (2 for the Euler-spiral model, 1 for the parabolic model), and N is sample size (=96, since each observer makes 48 settings of θ, and 48 settings of ϕ for each contour).[4] In the above expression, the first term captures the relative goodness of fit of the two models to the data, whereas the second term constitutes a penalty for more complex models (i.e., with more parameters). The fits to the parabolic model were determined using a likelihood function with the same functional form as Eq. (2), but with the "expected" angular positions $\theta_{p}(\kappa_0, r)$ and orientations $\phi_{p}(\kappa_0, r)$ at various radial distances r based on a parabola with curvature κ_0 at its vertex (the point of occlusion).

The Schwarz-criterion test showed that the Euler-spiral model was superior to the parabolic model in 31 of the 40 cases (see Table 4)—thereby indicating that visually extrapolated shape is *not* best characterized by a parabolic model (or, equivalently, a Taylor series truncated beyond the quadratic term). Specifically, in a large majority of the cases, the increase in the fit to the data was sufficiently high to warrant the additional parameter of the Euler-spiral model. Indeed, a comparison of the logarithmic-spiral model to the Euler-sprial model using the Schwarz criterion revealed the log-spiral model to be superior in 29 of the 40 cases.[5] The superior fits of the log-spiral model are consistent with our previous findings (Singh & Fulvio, 2005), and indicate that a nonlinear decrease in curvature, asymptoting to 0, better characterizes the shape of visually extrapolated contours than a linear decrease.

2.3.1. Scale invariance vs. short/long-range dichotomy

Our analysis has assumed that the same model of extrapolation applies for the entire range of distances tested, i.e., at all scales of interest—hence, scale invariance. Another possibility, however, is that there may be a principled distinction between 'short-range' and 'long-range' distances from the point of occlusion. Based in part on psychophyiscal results on the perception of contours in dot-sampled stimuli, Zucker and colleagues have proposed that two distinct stages are involved in the inference of contours from groups of discrete dots or line segments (see Link & Zucker, 1988; Zucker & Davis, 1988; Zucker, Dobbins, & Iverson, 1989). The first stage, operating over a relatively short range, includes receptive-field-based convolutions responsible for computing local orientation signals (represented coarsely at this stage), as well as local interactions between these convolutions, based on co-circularity. The second stage, operating over a larger range, is responsible for the inference of the global structure of the contour.

This account raises the possibility that different shape constraints may operate in contour extrapolation depending on whether they involve 'short range' or 'long range' interactions. It is possible, for instance, that the extrapolant shape may be based on co-circularity (i.e., constant-curvature extrapolation) at small distances from the point of occlusion, and the decaying curvature behavior obtains only at larger distances. Unfortunately, it is impossible to reliably distinguish between this hypothesis and a scale-invariant hypothesis from our data. If we restrict ourselves to the smallest measurement distance (0.68 deg)—or even the smallest two measurement distances (0.68 and 1.35 deg)—used in our experiment, the predictions of a circular-arc model and those of an Euler-spiral model are statistically indistinguishable for our stimuli. The largest differences in predicted angular position (obtained for the highest magnitude of rate of change of curvature; $|\gamma| = 0.0245\ \text{deg}^{-2}$) are 0.11° and 0.43° for the first two measurement distances (with the corresponding differences in predicted orientation being 0.33° and 1.3°, respectively). It is only by including the larger measurement distances, as

[4] The Schwarz criterion is closely related to the Bayes Information Criterion (BIC); in fact, it equals twice the difference between the respective BIC values of the two models (see, e.g., Kass & Raftery, 1995).

[5] A lograrithmic spiral is characterized by the intrinsic Cesàro equation $\kappa(s) = \frac{1}{a+bs}$, so its curvature κ varies inversely with its arc length s.

we did in the main analyses reported above, that the predictions of alternative models can be reliably distinguished in our data. Testing for the possible manifestation of a short-range vs. long-range distinction in the context of contour extrapolation will thus require future measurements that are aimed directly at distinguishing between a circular-arc model and an Euler-spiral model (or some other model that can exhibit a decrease in curvature) at very short distances from the point of occlusion.

2.4. Detection of variation in curvature

The above analyses show that observers did not use the rate of change of curvature in extrapolating contour shape. Moreover, their visually extrapolated contours were consistently characterized by decreasing curvature, irrespective of whether the inducer curvature was increasing or decreasing as it approached the occluder. Do these results simply reflect the fact that observers were unable to detect the variation in curvature along the inducing contours used (i.e., for the particular combinations of curvature, rate of change of curvature, and arc length of the inducing contours used in the experiment)?

In order to address this question, we obtained additional measurements on three naïve observers. We used the same 10 arcs of Euler spirals as in the main experiment (2 values of curvature × 5 values of rate of change of curvature), with the same fixed arc length of 4.56 deg. Unlike the main experiment, however, the stimulus did not include an adjustable line probe. Observers simply indicated whether the curvature of the inducing contour was increasing or decreasing as it approaches the occluder. The occluder size was not manipulated; its radius was fixed at 2.03 deg (half of the largest radius used in the main experiment).

Each observer participated in five sessions, with the first one being practice. The inducing contour was presented at a randomly assigned orientation, as in the main experiment. Two of the experimental sessions presented the contour as concave up, the other two as concave down, with their order counterbalanced. Each session contained 240 trials, with each trial requiring a 2AFC response of "increasing" or "decreasing" curvature. Each observer thus ran a total of 960 experimental trials.

Fig. 10 plots the data for the three observers. Each data point corresponds to the mean of 96 trials, for a given inducing contour (i.e., a particular combination of curvature and rate of change of curvature). Also shown on the plots are the fits of Weibull functions to the data (Wichmann & Hill, 2001). Although a small bias is apparent in some cases, it is evident that all observers are able to distinguish perceptually between increasing and decreasing curvature on the inducing contours. Thus the failure to use variation in curvature in extrapolating contours cannot be attributed to a failure to detect the variation in curvature in the inducing contours used. (Note, in particular, that observers' performance in detecting the *positive* rate of change of curvature is consistently high for the larger value of γ used. Hence, the observed property of decreasing curvature of visually extrapolated contours cannot be due to a misperception of the sign of variation in curvature of the inducing contours.)

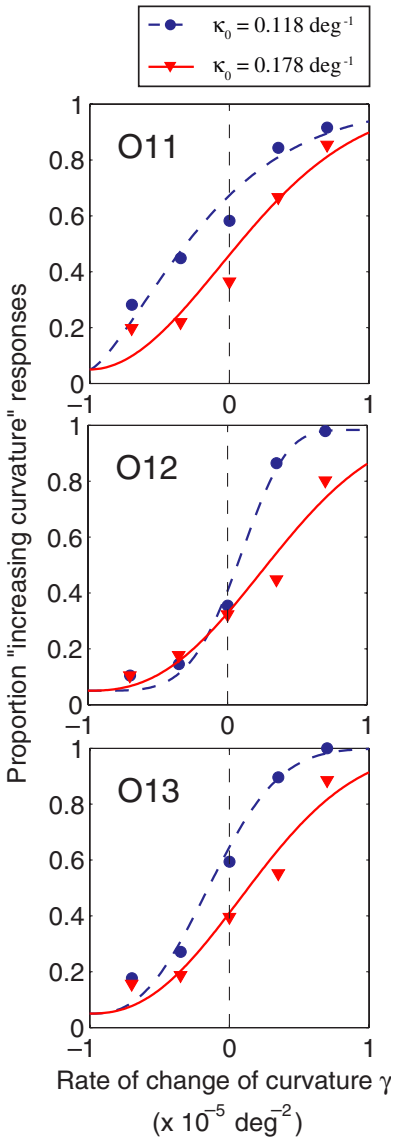

Fig. 10. Results in the task involving the detection of variation in curvature. The data indicate that observers can reliably identify whether the curvature of the inducing contour is increasing or decreasing; hence the failure to extrapolate rate of change of curvature does not simply reflect observers' inability to detect it.

3. Bayesian contour extrapolation

Our results indicate that human observers do not use the rate of change of curvature in extrapolating contour shape. Visually extrapolated contours are consistently characterized by decaying curvature, regardless of whether the inducing contours have increasing or decreasing curvature. Furthermore, extrapolation rate of change of curvature

exhibits no systematic dependence on the rate of change of curvature of the inducing contours. These results thus specify clear limits on the geometric properties that the visual system uses in "continuing" the shape of a contour: it uses tangent direction and curvature, but not rate of change of curvature.

These results also provide further support for a Bayesian model outlined by Singh and Fulvio (2005)—which is based on an interaction between the tendency to minimize curvature and the tendency to *continue* estimated curvature. The prior and the likelihood in the model are expressed as probability 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" (Elder & Goldberg, 2002; Feldman, 1997; Field et al., 1993; Geisler et al., 2001; Yuille et al., 2004), and is expressed as a Gaussian distribution on extrapolation curvature, centered on 0:

$$p(\kappa_{\text{ext}}) \sim N(0, \sigma_{\text{pr}}) \tag{5}$$

for some σ_{pr}. This bias is consistent in spirit with approaches that minimize total curvature along the length of the contour (Horn, 1983; Mumford, 1994), but is expressed as a probability distribution on local curvature. The likelihood bias captures the tendency toward co-circularity, namely, the tendency to "continue" the curvature of the inducing contour estimated at the point of occlusion (Elder & Goldberg, 2002; Feldman, 1997; Geisler et al., 2001; Parent & Zucker, 1989; Pizlo et al., 1997). It is consistent in spirit with approaches that minimize variation in curvature (Kimia et al., 2003), except that it takes into account the curvature of the inducing contour as well,[6] and is expressed as a distribution on local extrapolation curvature (centered on the estimated inducer curvature $\hat{\kappa}_i$ at the point of occlusion). 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: $\sigma_{\text{lik}}(d) = \sigma^0_{\text{lik}} + md$, where σ^0_{lik} is the standard deviation when the gap size is zero (infinitesimally thin occluder). The likelihood function is thus given by:

$$\ell(\kappa_{\text{ext}}|\hat{\kappa}_i, d) \sim N(\hat{\kappa}_i, \sigma^0_{\text{lik}} + md) \tag{6}$$

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

$$p(\kappa_{\text{ext}}|\hat{\kappa}_i, d) = \frac{\ell(\kappa_{\text{ext}}|\hat{\kappa}_i, d) \cdot p(\kappa_{\text{ext}})}{p(\hat{\kappa}_i)} \tag{7}$$

Under the assumption that the prior and likelihoods are both Gaussian distributions, there exist standard formulas for the posterior (Box & Tiao, 1992). In particular, the posterior is also a Gaussian with mean and variance given by:

$$\mu_{\text{post}}(\hat{\kappa}_i, d) = \left(\frac{\mu_{\text{pr}}}{\sigma^2_{\text{pr}}} + \frac{\mu_{\text{lik}}}{\sigma^2_{\text{lik}}(\hat{\kappa}_i, d)}\right) \Bigg/ \left(\frac{1}{\sigma^2_{\text{pr}}} + \frac{1}{\sigma^2_{\text{lik}}(\hat{\kappa}_i, d)}\right)$$

$$\sigma^2_{\text{post}}(\hat{\kappa}_i, d) = 1 \Bigg/ \left(\frac{1}{\sigma^2_{\text{pr}}} + \frac{1}{\sigma^2_{\text{lik}}(\hat{\kappa}_i, d)}\right) \tag{8}$$

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 the same expression for optimal combination (see Singh & Fulvio, 2006). 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., Clark & Yuille, 1990; Landy, Maloney, Johnston, & Young, 1995). Thus, consistent with the expression for the expected extrapolation curvature above:

$$\hat{\kappa}_{\text{ext}}(d) = w_{\text{pr}} \cdot \mu_{\text{pr}} + w_{\text{lik}} \cdot \mu_{\text{lik}} \tag{9}$$

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

Under the natural assumption that the continuation of estimated inducer curvature is subject to very little noise at the point of occlusion, we have $\sigma^0_{\text{lik}} \ll \sigma_{\text{pr}}$, and thus $w^0_{\text{lik}} \gg w_{\text{pr}}$. As a result, $\hat{\kappa}_{\text{ext}}(0) \approx \mu_{\text{lik}} = \hat{\kappa}_i$, 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—i.e., extrapolation curvature decreases asymptotically to zero. The rate of curvature decay is modulated by the slope term m:

$$\hat{\kappa}_{\text{ext}}(d) = \hat{\kappa}_i \cdot \frac{\sigma^2_{\text{pr}}}{\sigma^2_{\text{pr}} + (\sigma^0_{\text{lik}} + md)^2} \tag{10}$$

Fig. 11 plots this curvature decay for a number of different values of m.

This probabilistic cue-combination model thus explains the decaying-curvature behavior of visually extrapolated contours. Importantly, 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 extrapolation experiments with Euler spirals show, this property is consistent with human visual extrapolation.

[6] Consistent with the general calculus-of-variations approach to contour interpolation (e.g., Horn, 1983; Mumford, 1994), approaches based on minimizing variation in curvature (e.g., Kimia et al., 2003) obtain the total value of curvature variation by integrating only over the interpolated portion of the contour—not including the physically specified inducing contours. As a result, although the interpolated portion of the contour is second-order smooth, curvature discontinuities are generally introduced at points where the interpolated contour meets the inducing contours.

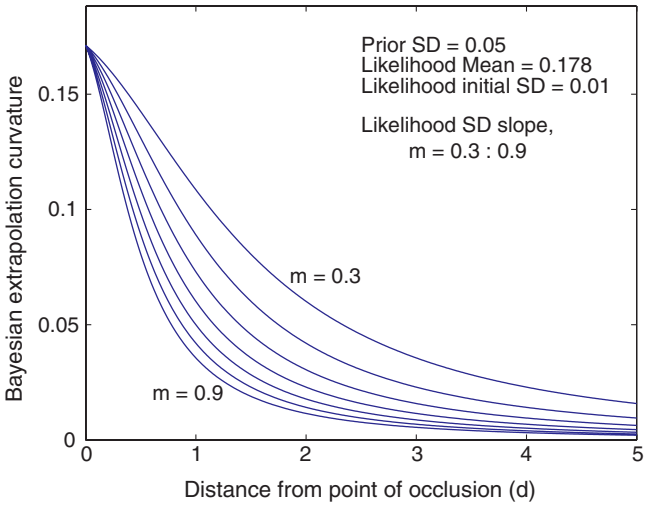

Fig. 11. Decay in the curvature of the Bayesian extrapolated contours, shown for different values of the slope term m.

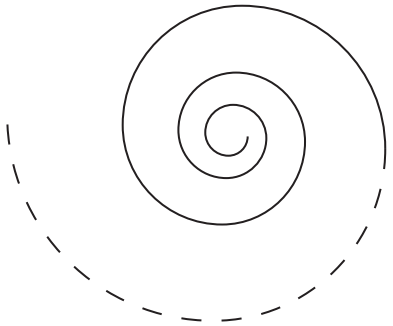

Fig. 12. An example of a spiral in which the "continuation" of the contour (dashed portion) is evident from other properties of the display—such as the systematic increase in radial distance from a clearly identifiable central region—rather than local estimates of rate of change of curvature.

4. Discussion

Our results specify clear limits on the order of geometric properties of the visible portion of a contour that the visual system uses in extending its shape. Specifically, the visual system makes systematic use of the tangent direction and curvature of a contour, but does not use its rate of change of curvature. The proposed Bayesian model specifies how these measurements on the inducing contour are used and combined to define the extended shape of a visually extrapolated contour. Under this model, the extrapolation shape results from an interaction between two constraints—minimization of variation in curvature (a "co-circularity" tendency to continue estimated curvature) and minimization of curvature (tendency toward "straightness")—with the relative weight assigned to the first constraint decreasing systematically with increasing distance from the point of occlusion.

The inability of observers to extrapolate the rate of change of curvature of a contour may seem surprising in light of examples of spiral displays, such as in Fig. 12, where the continuation of the spiral curve (marked with a dashed curve) seems straightforward. It should be noted, however, that in a display like this, there are other confounding factors that can support the estimates of continuation—most notably, the systematic increase in radial distance from a clearly identifiable central region (as well as, perhaps, the increasing separation between neighboring segments of the spiral, proceeding radially outwards). It thus seems likely that the extrapolation here is supported not by the rate of change in contour curvature per se, but by other such factors.[7] When these confounding factors are not readily available, such as in the stimulus configurations used in our experiments—so that the visual system is forced to rely on the local rate of curvature alone—we find a failure to make systematic use of this variable.

Our conclusion concerning the rate of change of curvature is consistent with that of a study by Feldman (1997). Observers in this study classified dot stimuli—generated by varying the turning angles between successive dots—as arising either from a curvilinear process, or independently. By analyzing observer responses ("curvilinear" vs. "random") using a probabilistic model of contour generation, Feldman found a substantial correlation between the influence of successive turning angles, but negligible correlation between that of non-successive turning angles. Thus, in addition to a prior for straightness (i.e, the distribution of turning angles is centered on 0, as embodied in the generative model), the visual system "expects" some covariance between successive turning angles. In other words, it expects that there will locally be only small changes in turning angle (the discrete form of curvature), as one proceeds along the contour. However, it does not similarly expect covariance between non-successive turning angles. The expectation of such covariance would require a bias toward small changes in *difference in successive turning angles* (the discrete version of rate of change of curvature). Based on this result, Feldman (1997) concluded that the visual system analyzes dot-sampled contours through local windows containing four successive dots (hence two turning angles) at a time. Thus, despite the differences between this study and ours—in terms of both stimuli (discrete dots vs. smooth contours) and the methodology (judgment of curvilinearity vs. measurement of extrapolation)—both studies point to a consistent conclusion: the visual system does not make systematic use of rate of change of contour curvature.

4.1. Probabilistic vs. variational approaches to shape completion

As noted earlier, variational approaches to shape completion assume a particular "energy" functional—such as

[7] Indeed, such factors may be used equally well to extend a spiral in the opposite—inward—direction, in which the curvature is increasing; despite the fact that both our present results and those from the previous study clearly indicate that visually extrapolated contours have systematically decreasing curvature.

total curvature (Horn, 1983; Mumford, 1994) or variation in curvature (Kimia et al., 2003)—to be minimized. This functional is, by definition, computed globally, i.e., by integrating over the length of an interpolated contour. As a result, variational approaches assume implicitly that the same shape constraint applies uniformly along the entire length of an interpolated contour. The analyses of our extrapolation measurements indicate, however, that the use of a fixed functional may not be appropriate for modeling shape completion by human vision. In particular, the relevant shape constraint itself can vary—e.g., the relative weights assigned to different shape constraints can vary—depending on distance from the point of occlusion. Hence, probabilistic models of shape completion, involving shape constraints expressed as probability distributions over local geometric variables (such as local curvature or orientation), are likely to provide a more appropriate class of models for human vision than variational models.

Probabilistic models are also richer than variational models. Apart from being able to model the variability present in human interpolation data, probabilistic models allow one to model systematic changes in the *shape* of interpolating contours, as a result of changes in the spread of the relevant probabilistic distributions. Consider, for instance, two inducing contours with the same curvature—one extremely short, and one long. The short contour is likely to generate a less reliable estimate of curvature than the long one. This increased variability in the case of the shorter inducing contour will translate into a quantifiable influence on interpolated shape—via probabilistic cue combination—a prediction that can then be compared against human visual interpolation.

Recent behavioral work has demonstrated the influence of factors beyond local contour geometry (i.e., beyond the relative positions and orientations of inducing edges) on the shape of visually interpolated contours. For instance, systematic shape differences have been shown to exist between partly occluded (i.e., amodal) and illusory (i.e., modal) contours (Anderson et al., 2002; Singh, 2004). Specifically, partly occluded contours are perceived as being more angular (i.e., closer to the intersection of the two individual extrapolants) than corresponding illusory contours. Within the context of probabilistic models of shape completion, this result is consistent with amodal completion having greater extrapolation strength (i.e., tighter spread) than modal completion (see Singh, 2004)—a prediction that can be independently tested. Similarly, in the context of stereoscopic illusory contours, the geometry of the *surface* enclosed by the contours has been shown to influence perceived illusory-contour shape. In particular, illusory contours bounding locally convex surfaces are perceived to be smoother than those bounding locally concave surfaces, and this influence of sign of curvature is modulated significantly by cross-axial shape width and medial-axis geometry (Fulvio & Singh, 2006). Again, probabilistic models permit sufficient flexibility to capture such shape differences, by manipulating the mean and variance (and possibly higher-order moments) of the relevant distributions, whereas it is more difficult to see how such influences could be accommodated within variational approaches.

4.2. Future directions

Our studies on visual extrapolation have so far used geometric contours—arcs of circles and parabolas in the previous study (Singh & Fulvio, 2005), and arcs of Euler spirals in the current study. The main benefit of using such contours is that they permit complete control over the curvature profile of the contours to be extrapolated—e.g., linearly increasing or decreasing curvature. These classes of curves were thus ideal, given our goal of characterizing "good continuation" in terms of testing (i) which local geometric properties of a contour the visual system uses in extrapolation, and (ii) how it combines these local properties to define extended shape of an extrapolated contour.

Having addressed these basic questions concerning the visual extrapolation of contours, however, a natural next step would be to examine visual extrapolation performance on more complex and naturalistic contours. Recent work on natural-image statistics has provided formal characterizations of contours found in images of the natural world (Elder & Goldberg, 2002; Geisler et al., 2001). In testing extrapolation performance on natural contours, one may thus either use one of these statistical models to generate a large number of contours, or use contours sampled directly from natural images. The primary question of interest would be how well the proposed Bayesian model for extrapolating smooth contours predicts performance on these more complex classes of contours, with structure at multiple scales. In other words, can the pattern of the visual system's sensitivities to contour properties, and the way in which it combines them to generate extrapolated shape, be understood in terms of the statistics of the natural contours? These and related important questions await further research.

Acknowledgments

Portions of this work were presented at the 5th Workshop on Perceptual Organization in Computer Vision at the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2006 (see Singh & Fulvio, 2006). We are grateful to Jacob Feldman, Randy Gallistel, Eileen Kowler, and Larry Maloney, as well as two reviewers, for helpful comments and suggestions. This work was supported by NSF grant BCS-0216944.

References

Anderson, B. L., & Barth, H. C. (1999). Motion-based mechanisms of illusory contour synthesis. *Neuron, 24*, 433–441.

Anderson, B. L., Singh, M., & Fleming, R. (2002). The interpolation of object and surface structure. *Cognitive Psychology, 44*, 148–190.

August, J., & Zucker, S. W. (2001). A Markov process using curvature for filtering curve images. In M. A. T. Figueredo, J. Zeurbia, & A. K. Jain (Eds.), *EMMCVPR 2001, energy minimization methods in computer vision and pattern recognition* (pp. 497–512). Sophia Antipolis, France: Springer-Verlag.

Barrow, H. G., & Tenenbaum, J. (1981). Interpreting line drawings as three-dimensional surfaces. *Artificial Intelligence, 17*, 75–116.

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

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

Box, G. E. P., & Tiao, G. C. (1992). *Bayesian inference in statistical analysis*. New York: Wiley.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10*, 433–436.

Bregman, A. L. (1981). Asking the "what for" question in auditory perception. In M. L. Kubovy & J. R. Pomerantz (Eds.), *Perceptual organization* (pp. 141–180). Hillsdale, NJ: Lawrence Erlbaum.

Caelli, T. M., & Umansky, J. (1976). Interpolation in the visual system. *Vision Research, 16*, 1055–1060.

Clark, J. J., & Yuille, A. L. (1990). *Data fusion for sensory information processing systems*. Boston, MA: Kluwer.

Dobbins, A., Zucker, S. W., & Cynader, M. S. (1987). Endstopped neurons in the visual cortex as a substrate for calculating curvature. *Nature, 329*(6138), 438–441.

Elder, J. H., & Goldberg, R. M. (2002). Ecological statistics of Gestalt laws for the perceptual organization of contours. *Journal of Vision, 2*(4), 324–353.

Euler, L. (1744/1952). De curvis elasticis. In C. Carathéodory (Ed.), *Methodus inventiendi lineas curvas maximi minimive proprietate gaudentes sive solutio problematis isoperimetrici latissimo sensu accepti. Opera Omnia, Ser. I.: Opera Mathematica* (Vol. 24, pp. 308–342). Lausanne: Birkhauser [Originally published in 1744 in book format].

Fantoni, C., & Gerbino, W. (2003). Contour interpolation by vector-field combination. *Journal of Vision, 3*(4), 281–303.

Feldman, J. (1997). Curvilinearity, covariance, and regularity in perceptual groups. *Vision Research, 37*(20), 2835–2848.

Feldman, J. (2001). Bayesian contour integration. *Perception & Psychophysics, 63*(7), 1171–1182.

Feldman, J., & Singh, M. (2005). Information along contours and object boundaries. *Psychological Review, 112*(1), 243–252.

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

Fulvio, J., & Singh, M. (2006). Surface geometry influences the shape of illusory contours. *Acta Psychologica, 123*, 20–40.

Geisler, W. S., Perry, J. S., Super, B. J., & Gallogly, D. P. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research, 41*(6), 711–724.

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

Guttman, S. E., & Kellman, P. J. (2004). Contour interpolation revealed by a dot localization paradigm. *Vision Research, 44*, 1799–1815.

Heitger, F., von der Heydt, R., Peterhans, E., Rosenthaler, L., & Kübler, O. (1998). Simulation of neural contour mechanisms:representing anomalous contours. *Image and Vision Computing, 16*, 407–421.

Hon, A., Maloney, L., & Landy, M. (1997). The influence function for visual interpolation. *SPIE Conference Proceedings, 3016*, 409–419.

Horn, B. K. P. (1983). The curve of least energy. *ACM Transactions on Mathematical Software, 9*(4), 441–460.

Kanizsa, G. (1979). *Organization in vision: Essays on Gestalt perception*. New York: Präger.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association, 90*(430), 773–795.

Kellman, P. J., & Shipley, T. F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology, 23*, 141–221.

Kimia, B. B., Frankel, I., & Popescu, A. (2003). Euler spiral for shape completion. *International Journal of Computer Vision, 54*(1/2), 157–180.

Klein, S. A., & Levi, D. M. (1987). Position sense of the peripheral retina. *Journal of the Optical Society of America, A, 4*, 1543–1553.

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.

Kubovy, M., & Gepshtein, S. (2000). Optimal curvatures in the completion of visual contours. *Investigative Ophthalmology and Visual Science, 41*(Suppl.), B568.

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

Link, N., & Zucker, S. W. (1988). Corner detection in curvilinear dot grouping. *Biological Cybernetics, 59*, 247–256.

Love, A. E. H. (1927). *Mathematical theory of elasticity*. Cambridge: University Press.

Mardia, K. V., & Jupp, P. (2000). *Directional statistics*. New York: Wiley Series in Probability and Statistics.

Mood, A., Graybill, F. A., & Boes, D. C. (1974). *Introduction to the theory of statistics* (3rd ed.). New York: McGraw-Hill.

Mumford, D. (1994). Elastica and computer vision. In C. L. Bajaj (Ed.), *Algebraic geometry and its applications* (pp. 491–506). New York: Springer-Verlag.

Nakayama, K., He, Z., & Shimojo, S. (1995). Visual surface representation: A critical link between lower-level and higher-level vision (2nd ed.. In S. M. Kosslyn & D. N. Osherson (Eds.). *Visual cognition: An invitation to cognitive science* (Vol. 2, pp. 491–506). Cambridge, MA: MIT Press.

Parent, P., & Zucker, S. W. (1989). Trace inference, curvature consistency and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 11*, 823–839.

Pelli, D. G. (1997). The Videotoolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*, 437–442.

Pettet, M. W., McKee, S. P., & Grzywacz, N. M. (1998). Constraints on long range interactions mediating contour detection. *Vision Research, 38*(6), 865–879.

Pizlo, Z., Salach-Goyska, M., & Rosenfeld, A. (1997). Curve detection in a noisy image. *Vision Research, 37*(9), 1217–1241.

Ringach, D. L., & Shapley, R. (1996). Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision Research, 36*(20), 3037–3050.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics, 6*, 461–464.

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

Singh, M. (2004). Modal and amodal completion generate different shapes. *Psychological Science, 15*, 454–459.

Singh, M., & Fulvio, J. M. (2005). Visual extrapolation of contour geometry. *Proceedings of the National Academy of Sciences, USA, 102*(3), 939–944.

Singh, M., & Fulvio, J. (2006). Contour extrapolation using probabilistic cue combination. In Proceedings of the conference on computer vision and pattern recognition (Workshop on perceptual organization in computer vision). http://dx.doi.org/10.1109/CVPRW.2006.61.

Singh, M., & Hoffman, D. D. (1999). Completing visual contours: The relationship between relatability and minimizing inflections. *Perception & Psychophysics, 61*, 636–660.

Smits, J. T., & Vos, P. G. (1987). The perception of continuous curves in dot stimuli. *Perception, 16*, 121–131.

Takeichi, H. (1995). The effect of curvature on visual interpolation. *Perception, 24*, 1011–1020.

Takeichi, H., Nakazawa, H., Murakami, I., & Shimojo, S. (1995). The theory of the curvature-constraint line for amodal completion. *Perception, 24*, 373–389.

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

Uttal, W. R. (1973). The effect of deviations from linearity on the detection of dotted line patterns. *Vision Research, 13*, 2155–2163.

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

Warren, P. A., Maloney, L. T., & Landy, M. S. (2004). Interpolating sampled contours in 3D: Perturbation analyses. *Vision Research, 44*(8), 815–832.

Wichmann, F. A., & Hill, N. J. (2001). The psychometric function: I. Fitting, sampling and goodness-of-fit. *Perception & Psychophysics, 63*(8), 1293–1313.

Yuille, A. L., Fang, F., Schrater, P., & Kersten, D. (2004). Human and ideal observers for detecting image curves. In S. Thrun, L. Saul, & B. Schölkopf (Eds.). *Advances in neural information processing systems* (Vol. 16, pp. 59–70). Cambridge, MA: MIT Press.

Zucker, S. W., & Davis, S. (1988). Points and endpoints: A size/spacing constraint in dot grouping. *Perception, 17*, 229–247.

Zucker, S. W., Dobbins, A., & Iverson, L. (1989). Two stages of curve detection suggest two styles of visual computation. *Neural Computation, 1*, 68–81.