

Precision and consistency of contour interpolation

Jacqueline M. Fulvio^{a,*}, Manish Singh^b, Laurence T. Maloney^{a,c}

^a Department of Psychology, New York University, 6 Washington Place Rm. 877C, New York, NY 10003, USA

^b Department of Psychology and Center for Cognitive Science, Rutgers University, New Brunswick, USA

^c Center for Neural Science, New York University, USA

Received 2 January 2007; received in revised form 20 December 2007

Abstract

We investigated conditions under which observers can interpolate occluded contours by a single, stable, smooth contour. Observers viewed partly-occluded contours defined by linear segments and estimated the position and tangent orientation of the contour at multiple locations within the occluded region. We measured the precision and consistency of observers' settings as indices of successful interpolation. We found that although increasing the relative offset between inducers led to a decrease in both precision and consistency, increasing turning angle affected primarily precision. We discuss conditions under which interpolation settings are consistent with a single, stable smooth contour.

© 2008 Elsevier Ltd. All rights reserved.

Keywords: Contour interpolation; Visual completion; Reliability; Consistency

1. Introduction

The problem of perceptual organization was first formulated by Gestalt psychologists. They attributed special “emergent” perceptual properties to whole objects—properties that cannot be achieved by “piecing” together image fragments. The two displays in Fig. 1 illustrate one such Gestalt phenomenon: the amodal completion of partly-occluded contours. The first display (Fig. 1A) is readily interpreted as an occluded contour; the two line segments appear to belong to a single continuous contour that extends behind the occluder. In the second display (Fig. 1B), however, the two line segments are perceived as disparate and unconnected—not belonging to a single extended contour.

Previous work has shown that mechanisms of amodal completion do in fact generate a representation that goes beyond the fragmented image data—with measurable consequences for perception and attention (e.g., Anderson, Singh, & Fleming, 2002; Davis & Driver, 1998; Guttman,

Sekuler, & Kellman, 2003; Liu, Jacobs, & Basri, 1999; Rauschenberger & Yantis, 2001; Ringach & Shapley, 1996; Sekuler & Palmer, 1992). Mechanisms of amodal completion allow for a relatively stable representation of a scene in the face of continual changes in the observer's vantage point, which lead to different portions of the scene becoming occluded or disoccluded. Moreover, there are contexts such as manual grasping of a partly-occluded object, where an accurate “completed” representation is critical. Very little past work has, however, obtained detailed measurements of amodally interpolated shape and its dependence on the geometry of the two inducing edges.

The Gestalt psychologists accounted for the difference in interpretation between the two displays in Fig. 1 by the principle of *good continuation*: the line segments in Fig. 1A have “good continuation” and can be grouped into a single contour (Michotte, Thinès, & Crabbé, 1964; Wertheimer, 1923). Those in Fig. 1B do not, and are thus unlikely to be grouped together. In its original formulation, however, this principle is vague and lacks predictive power. This has motivated recent attempts to characterize good continuation more formally, in terms of geometric relations

* Corresponding author.

E-mail address: jmf384@nyu.edu (J.M. Fulvio).




Fig. 1. Two stimulus configurations illustrating ‘good continuation’ and its absence. (A) The two contour fragments satisfy Kellman and Shipley’s (1991) reliability criteria and, according to Kellman & Shipley’s, should be perceived as part of a single smooth contour. (B) The two contour fragments violate these reliability criteria and Kellman & Shipley predict that the observer will not perceive them as part of a single smooth contour.

that support contour grouping (Elder & Goldberg, 2002; Field, Hayes, & Hess, 1993; Geisler, Perry, Super, & Gallogly, 2001; Grossberg & Mingolla, 1985; Kellman & Shipley, 1991; Parent & Zucker, 1989).

Characterizations of good continuation have been proposed in two closely related domains: visual completion of partly occluded and illusory contours, and visual integration of discrete oriented elements into smooth, extended contours. Research in the first domain typically focuses on stimuli consisting of two segments of smooth inducing contours separated by an intervening gap (representing partial occlusion or the effect of camouflage) that must be interpolated (see Fig. 1A). In this domain, Kellman and Shipley (1991) proposed the reliability hypothesis, according to which the visual system will interpolate a contour between two inducing contours if their edges (at their respective points of occlusion) satisfy two criteria: (i) the linear extensions of the two edges intersect, and (ii) their turning angle (the angle through which one edge must turn to align with the other) does not exceed 90° . We will refer to these two criteria as the *linear extension intersection criterion*, and the *turning angle cutoff criterion*, respectively. The two line segments in Fig. 1A satisfy both reliability criteria and, according to the reliability hypothesis, they can be visually grouped and completed; whereas those in Fig. 1B violate both. The reliability of two edges has been shown to influence ‘completion strength’, based on observer ratings (Kellman & Shipley, 1991; Shipley & Kellman, 1992), depth discrimination (Liu et al., 1999; Yin, Kellman, & Shipley, 2000), and shape (fat/thin) discrimination (Ringach & Shapley, 1996).

The related domain of contour integration investigates conditions under which a set of discrete elements, such as short line segments or Gabor patches, is integrated into the representation of a single extended contour (Fig. 1B). Integration performance is found to be best when pairs of successive elements are cocircular, i.e., tangent to a common circle (Parent & Zucker, 1989), and decays with deviation from cocircularity. In addition, there is a strong influence of turning angle: grouping is strongest when suc-

cessive pairs of edges are collinear, and decreases systematically with increasing turning angle between them (Elder & Goldberg, 2002; Feldman, 1997, 2001; Field et al., 1993; Geisler et al., 2001; Pizlo, Salach-Golyska, & Rosenfeld, 1997). These dependencies are described in terms of an ‘association-field’ model (Ben-Shahar & Zucker, 2004; Field et al., 1993; Geisler et al., 2001), which summarizes the grouping strength between pairs of oriented elements, as a function of their relative positions and orientations. The pattern of perceptual grouping strength is found to be consistent with the statistics of edge pairs found in natural images, which also exhibit a dominance of cocircular and collinear structure—thereby supporting the idea that over the course of evolution the brain has internalized regularities found in the natural environment (Elder & Goldberg, 2002; Geisler et al., 2001; Sigman, Guillermo, Gilbert, & Magnasco, 2001).

What is the relationship between contour reliability and the constraints based on cocircularity and collinearity embodied in the association-field geometry? Both reliability and association-field geometry of course involve a dependence on turning angle. Beyond this, however, it is easily seen that the *linear extension intersection criterion* of reliability is weaker than cocircularity. A pair of cocircular edges necessarily satisfies the *extension criterion* of reliability, but not vice versa. The *extension criterion* is formally equivalent to the existence of a non-inflecting smooth contour that interpolates between the two edges (Singh & Hoffman, 1999). Hence, whereas cocircularity requires a smooth interpolating contour of constant curvature between two edges, the *extension criterion* requires only a smooth interpolating contour of uniform *sign* of curvature.

Another point of difference, however, is that the reliability hypothesis articulates both of its criteria in terms of binary properties: either the turning angle is less than 90° or it is not, either the linear extension intersection criterion is satisfied or not, and either interpolation is successful—or it is not. Thus in its original formulation, it predicts abrupt, qualitative changes in interpolation performance as a result of smooth, continuous changes in the stimulus configuration. If we were to vary the turning angle in Fig. 1A, for example, gradually increasing it to 90° and beyond, the reliability hypothesis entails that performance will abruptly change at 90° . Similarly, as one gradually increases the vertical offset between the two inducers in Fig. 1A, it entails that performance will change abruptly at some point (when their linear extensions fail to intersect). Such binary or categorical shifts in interpolation performance are implausible. Indeed, Kellman and Shipley themselves appear to embrace graded versions of their criteria (Kellman & Shipley, 1991; Kellman, Guttman, & Wickens, 2001), notwithstanding the formal articulation of these criteria. The association-field type model, on the other hand, explicitly embodies a graded, probabilistic dependence of grouping on geometric variables.

In what follows we will treat both reliability criteria as continuous. Our experiments will examine how interpolation

performance varies with two geometric variables embodied in the relatability criteria and in the association-field model: the turning angle between the two inducers, and their relative “vertical” offset (which determines how far their extensions are from intersecting, as well as how far they are from being cocircular).

The *relatability hypothesis* and the association field model both concern *grouping*—how likely two edges are to be perceptually grouped into the representation of a single extended contour. They make no explicit predictions concerning the shape of a visually completed contour. A handful of previous studies have examined the shapes of visually interpolated contours. These studies have made use of a variety of dependent measures, intended to characterize shape. Such measures have included reports of the perceived number of inflections along partly-occluded contours (Takeichi, Nakazawa, Murakami, & Shimojo, 1995); localization of extremal points—i.e., points farthest from the points of occlusion—along the interpolated contour (Fantoni & Gerbino, 2003; Takeichi, 1995); adjustment of points to lie on the missing portion of a contour in two-dimensional space (Anderson & Barth, 1999; Guttman & Kellman, 2004; Hon, Maloney, & Landy, 1997) and three-dimensional space (Warren, Maloney, & Landy, 2002, 2004); and parametric shape matching (Fulvio & Singh, 2006; Singh, 2004). These studies have provided important information on the influence of various geometric variables on visually interpolated shape. However, few previous studies have mapped out the *extended* shapes of visually interpolated contours (but see Anderson & Barth, 1999, who obtained positional measurements at multiple locations along motion-defined illusory contours). And few studies measuring interpolation shape have employed configurations such as Fig. 1B where grouping and therefore interpolation are unlikely to occur.

The following series of experiments investigates contour interpolation—i.e., observers’ ability to localize a partly-occluded contour—under a variety of geometric conditions where the two relatability criteria are violated to a lesser or greater extent. One evident difficulty in testing the relatability hypothesis is that it predicts conditions under which contour interpolation does not occur or is in some way impaired. We therefore begin by formulating two explicit measures of the degree to which successful grouping and interpolation has occurred.

The first measure is based on the precision of repeated settings. Intuitively, conditions that are highly conducive for visual interpolation should evoke precise settings from trial to trial (i.e., low setting variability). Conversely, under image conditions that do not support visual completion, no smooth contour should provide a “good” interpolation between the two inducers; hence observers’ repeated settings should be highly variable. A number of researchers have sought to measure completion “strength” using different methods (e.g., Anderson et al., 2002; Kellman & Shipley, 1991; Liu et al., 1999; Ringach & Shapley, 1996; Takeichi et al., 1995). It is plausible that reduced precision

in interpolation settings would correspond to greater subjective uncertainty in localizing the contour. It would not be surprising if it correlated with contour “strength” as proposed by previous studies. We take as our measure of imprecision the standard deviation of observer’s settings.

The second measure we propose is based not on the variability of a repeated setting but on the self-consistency of the observer’s multiple judgments concerning a single contour. In order to measure consistency, we ask observers to make both position and tangent orientation judgments at several locations along the contour in the region of occlusion (see, e.g., Singh & Fulvio, 2005, 2007 who obtained paired measurements of position and tangent orientation at multiple distances from the point of occlusion in the context of contour extrapolation). As described below, these position and orientation judgments can be compared to test the internal consistency of observers’ performance—to determine whether their settings are consistent with any single contour. Such tests of consistency are based on comparisons of settings made by the observer at multiple locations. Settings at any single location cannot, by themselves, lead to rejection of consistency. We use two measures of inconsistency, one based on polynomial fits, and one non-parametric.

It is generally believed that, given a fixed stimulus, visual mechanisms of contour interpolation lead to the formation of a single, stable, smooth interpolating contour. This assumption remains untested, however. Our consistency measure will allow us to test this *single, stable smooth contour hypothesis* and, as we shall see, reject it in many cases. We abbreviate “*single, stable, smooth*” by S^3 for convenience.

The two kinds of failure—based on our two measures of interpolation performance—need not be mutually exclusive. Observers’ settings can be both less precise and less consistent in some experimental conditions than they are in others. We will look for evidence of both kinds of failure in investigating how inducer geometry influences contour interpolation. We will test whether relatability as defined by either of its two criteria permits precise, consistent contour interpolation, and whether a decrease in relatability by either criterion leads to a decrease in precision, a decrease in consistency, both, or neither.

2. General methods

2.1. Observers

Three observers participated in the experiments at Rutgers University. Two were not aware of the purpose of the experiments and one was an author (JMF). All had normal or corrected-to-normal vision.

2.2. Stimuli

Examples of the stimulus display are shown in Fig. 2. Each display contained two oriented line segments referred to as the *inducers*. The inducers were white, 3.34 degrees of visual angle (DVA) long, and 0.028 DVA thick, with anti-aliasing at the resolution of one-fourth of a pixel. They varied with respect to their *orientation*—the angle θ ($0 < \theta < \frac{\pi}{2}$) of

the inducer from the horizontal (see Fig. 3A). Note that we measure θ for the left inducer in the counterclockwise direction, and for the right inducer in the clockwise direction. The two angles labeled θ in the left diagram in Fig. 3A are equal and both positive whereas in the right diagram they are not equal on account of the skew (γ) manipulation described below in Experiment 1b.

We varied two independent measures of interest: (i) the *turning angle* τ between a pair of inducers with orientations θ_1, θ_2 , which is equal to the sum of their orientations, $\tau = \theta_1 + \theta_2$; and (ii) the relative vertical offset

between each pair of inducers at their respective points of occlusion, denoted Δ . Bearing in mind our continuous interpretation of Kellman and Shipley's (1991) turning angle criterion (see above), any inducer pair is non-relatable to the extent that its turning angle τ exceeds 90° . Furthermore, by the inducer extension intersection criterion, an inducer pair is relatable if and only if

$$|\Delta| < w \tan \theta \tag{1}$$




Fig. 2. Examples of stimulus displays used in the studies. The observer could control the height and tangent orientation of the short linear segments that appear in the vertical slit in the occluded region.

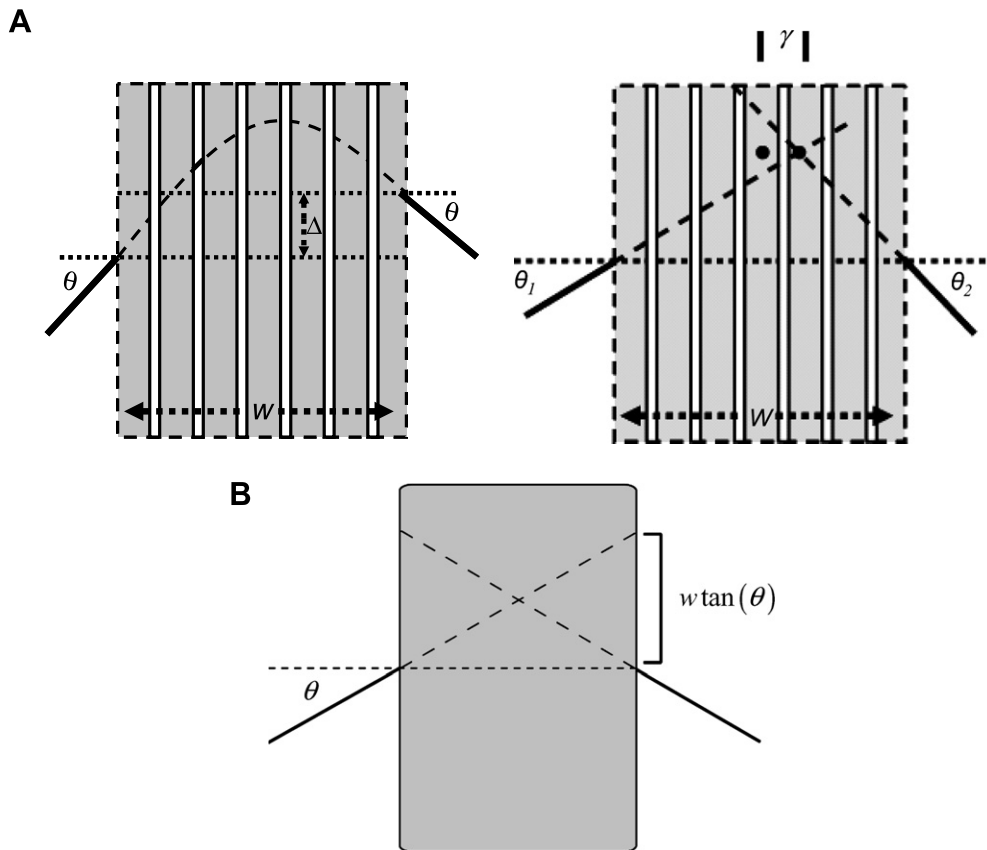


Fig. 3. Schematic of the displays used in the studies. (A) All six of the vertical interpolations windows are shown. Only one was visible on each trial. There were three manipulations: the orientation of the inducers (denoted by θ), the relative vertical offset between the inducers (denoted by Δ) and the amount of offset of the inducer extension intersection point relative to the horizontal midpoint of the occluder ("skew" denoted by γ). See text for details. (B) Illustration of the inducer extension intersection criterion. Inducers that are vertically offset by more than $w \tan \theta$ are non-relatable.

as illustrated in Fig. 3B. To the extent that $|\Delta|$ exceeds $w \tan \theta$ the inducers are non-relatable. The specific values of turning angle and Δ will be provided in the Methods sections for each of the three experiments.

In each experiment, the inducers pointed upward in half of the experimental sessions, and downward in the other half. One of the two inducers was designated the reference inducer, and its vertical height was used as a reference for the vertical placement of the opposite inducer (i.e., along the opposite vertical edge of the occluder). On any given trial, the reference inducer could appear either on the left or on the right side.

The inducers were placed at the left and right edges of an occlusion region shown as a gray rectangle with rounded corners, having width w equal to 3.5 DVA and height equal to 9.10 DVA. On each trial, a narrow vertical slit of width 0.14 DVA, appeared within the occluder at one of six horizontal locations, which we refer to as an *interpolation window*. (To aid in specifying the stimuli, all six windows are shown in the schematics in Fig. 3A, even though only one would appear on any given trial.) The vertical midlines of the windows were 0.42 DVA apart, and the leftmost and rightmost windows were 0.49 DVA from the closest vertical edge of the occluder. None of the windows appeared at the horizontal midpoint of the occluder.

Through the interpolation window, a white, straight-line probe was visible, whose vertical position and orientation were to be adjusted by the observer. The probe had the same color, thickness, and anti-aliasing as the inducing contours. It was initially presented with a horizontal orientation, at the vertical midpoint of the interpolation window. In the position-adjustment mode, the probe moved vertically within the window; in the orientation-adjustment mode, it pivoted about its midpoint (which was constrained to lie on the vertical midline of the window). The position or “height” of the probe, denoted h , was constrained to lie between the uppermost and lowermost horizontal edges of the occluder. It was measured relative to the height of the reference inducer at its point of occlusion. The orientation of the probe, denoted ϕ , could range from -90° to $+90^\circ$, which allowed for the full range of orientations.

2.3. Software and apparatus

The stimulus displays were presented on a high-resolution 22-in. monitor (*Lacie Blue*) with a display area of 40.3 cm \times 30.2 cm and, a resolution of 1600 \times 1200 pixels in conditions of low ambient illumination. Observers viewed the stimuli from a distance of 102.5 cm, their viewing location fixed by means of a head and chin rest. All stimuli were displayed using the Psychophysics Toolbox extensions for MATLAB (Brainard, 1997; Pelli, 1997) and were presented in the center of the screen on a black background.

2.4. Procedure

The task of the observers in all three experiments was to adjust the position, h , and orientation, ϕ , of the straight-line probe until the combination of settings optimized the percept of a smoothly-continuing partly-occluded contour defined by the two inducing contours.

Within each experiment, the observers participated in four experimental sessions, preceded by a practice session. Within a session, each pair of inducers (thus, each combination of turning angle, τ , and relative inducer offset, Δ) was presented twice for each of the six window locations (with the reference inducer presented once on the left and once on the right).

In half of the experimental sessions, we presented the inducing contours pointing upward and the other half presented them pointing downward. The order was counterbalanced. On any given trial, observers first adjusted the position of the line probe vertically within the window, using a trackball. Pressing the space bar then allowed them to toggle to adjusting the orientation of the line probe, while maintaining its height setting. Observers toggled back and forth in this manner between height and tangent orientation settings in order to optimize the percept of a smooth partly-occluded contour defined by the two inducing contours. They pressed the trackball button when they were satisfied with the combination of height and tangent orientation settings. A question appeared at the bot-

tom 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.

3. Experiment 1a

We begin our investigation by testing the effect of turning-angle on contour interpolation. Recall that the postulated cutoff for visual interpolation is 90° , beyond which visual interpolation should degrade. We chose six turning angles, three less than 90° and three greater. All other aspects of the stimuli were equated—length, separation, and relative height of the inducers. If visual interpolation does in fact degrade to the extent that the turning angle between inducers exceeds 90° , observers’ setting precision should decrease with increasing turning angle, indicating that the visually interpolated contours are less well localized, with a categorical shift occurring at 90° . Moreover, consistency between observers’ position and tangent orientation settings should also deteriorate.

3.1. Methods

3.1.1. Inducers

The turning angle (τ) between the inducer pairs took on one of six values: 40° , 60° , 80° , 100° , 120° , 140° (which correspond to individual inducer orientations (θ) of 20° , 30° , 40° , 50° , 60° , 70° relative to the horizontal, respectively). Thus, the latter three turning angles are non-relatable. No vertical offset between inducers was applied, so all inducer pairs had $\Delta = 0$. We note that the inducer pairs in this experiment were also co-circular (the two inducers were tangent to a single circle at their respective points of occlusion).

3.1.2. Design

The design contained six turning angles and six window locations. The reference occluder could appear either on the left or the right side. Given the symmetry of the stimuli (since $\Delta = 0$), this last factor could have no effect in this experiment; but it will become relevant in later experiments. Each session thus contained $6 \times 2 \times 6$ or 72 trials, with each trial obtaining paired settings of position and tangent orientation of the line probe. Each observer performed adjustments in 4 experimental sessions.

3.2. Results and analysis

On each trial we recorded the observer’s position and tangent orientation settings. For each inducer-pair geometry, measurements were taken through each of the six windows. Following preliminary analysis, we were able to combine conditions that differed only in left-right reflection and upwards-downwards inducer presentation. We will describe the stimuli and settings as if all stimuli were transformed to common coordinates in all three experiments. For each of the remaining six conditions (6 values of

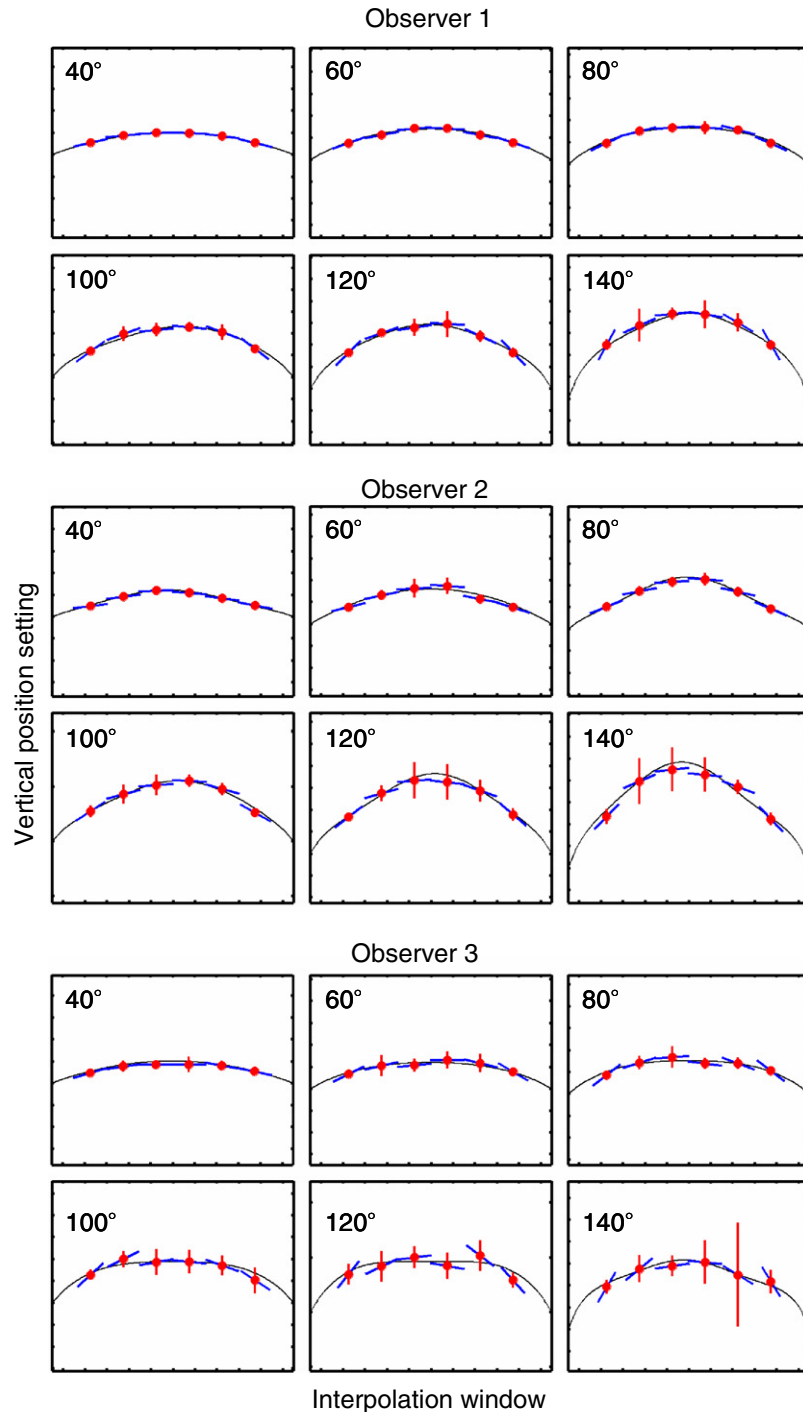


Fig. 4. Experiment 1a: Mean settings and polynomial fits. The means and standard deviations of position settings are shown as points with error bars (± 1 SD). The means of the tangent orientation settings are represented by short line segments whose tangent orientation is the mean setting. The superimposed curves are the best-fitting polynomials. Details concerning polynomial fits will be provided below.

turning angle τ), we computed the mean position settings h_1, \dots, h_6 , and mean tangent orientation settings $\varphi_1, \dots, \varphi_6$, for each of the interpolation windows. Fig. 4 shows the interpolation data for each of the three observers, in each of the six inducer-pair conditions. The curves in these graphs were fit to the data as described in the Section 3.2.2 below.

3.2.1. Analysis of precision

To analyze how the precision of an interpolating contour is affected by the turning angle between inducers, we computed the standard deviations of the position and tangent orientation settings for each condition. (This was done by first computing the standard deviations of the eight position settings within each interpolation window, and

then collapsing these standard deviations across all six windows using root mean squares.) This procedure was repeated for the tangent orientation settings. Fig. 5 shows the standard deviations plotted as a function of turning angle between inducers for all three observers with 95% confidence intervals, for the two settings types. As expected, standard deviations increase with increasing turning angle for both position and tangent orientation settings, for all three observers—thereby indicating that the interpolating contours become less sharply defined with increasing inducer turning angle. However, in its original discrete formulation, the reliability hypothesis predicts that once the 90° cutoff has been exceeded, visual completion should no longer occur, and hence precision should decline drastically. Our results indicate that this is not the case, but rather there is a *gradual*, roughly linear increase in variability (or imprecision) with turning angle—both reliable and non-reliable—with no evidence of a sudden increase beyond the postulated cutoff point of 90°.

3.2.2. Analysis of shape and consistency

3.2.2.1. Shape. To characterize the shape of the visually interpolated contours implied by observers’ settings, we first fit polynomials to the *position* data only. During fitting, all polynomials were forced to match the slope and position of the two inducers at their respective points of occlusion (four constraints). An *n*th degree polynomial has *n* + 1 free parameters. When the inducers are symmetrically placed, as in this experiment, a quadratic polynomial (*n* = 2) satisfies these inducer constraints. As a result, we are matching the data to an otherwise parameter-free fit with this polynomial (i.e., the inducers define a unique parabola). When $\Delta \neq 0$, i.e., when the inducers are vertically offset from each other, polynomials of at least degree *n* = 3 are needed to satisfy the inducer constraints. (This will become important in Experiment 2.) Beyond that, when *n* > 3, the imposed constraints do not uniquely

determine the resulting polynomial, but instead specify a family of polynomials. For those cases, we indexed the resulting families of polynomial curves by adding *n* – 3 additional position constraints at control points whose horizontal locations are spaced equally across the occlusion region. The vertical positions of these control points function as free parameters that we could adjust in matching polynomials of a specified degree to observers’ position settings. For *n* ≤ 3, we calculated the sum of squared differences (SS) between the unique polynomial and the observer’s position settings. For *n* > 3, we numerically varied control points to select the single polynomial in the family that minimized the SS.

For each condition, we obtained maximum likelihood estimates (Mood, Graybill, & Boes, 1974, pp. 276ff) for the fits of polynomials of degrees 2–9 (recall that the parabola, i.e., polynomial of degree 2, is uniquely defined by the inducers when $\Delta = 0$). We record the log likelihood associated with each fit ($\lambda_n, n = 2, 3, \dots, 9$). We performed nested hypothesis tests (Mood et al., 1974, p. 441ff) to determine whether the addition of more parameters yields a significantly better fit. This was done by computing the log likelihood ratios between two polynomial fits using the following formula:

$$H = 2 * (\lambda_{n+m} - \lambda_n) \tag{2}$$

and comparing them to the χ^2_m distribution, where the degrees of freedom *m* corresponds to the difference in the number of parameters of the two fitted polynomials. We rejected the null hypothesis when $H > \chi^2_m(0.01)$, the cutoff for a hypothesis test of size (critical value) 0.01. The tests were performed in a hierarchical order, as depicted in Fig. 6A. The log likelihood value will never decrease as we increase the degree of the fitting polynomial since a lower degree polynomial is a special case of any higher degree polynomial. The significance tests in effect measure whether the improvement in fit (increased log likelihood) obtained by increasing the degree of the polynomial is large enough to justify selection of the higher degree polynomial.

We note that our choice of parameterization of the polynomials (by control points) cannot affect the outcome of the maximum likelihood estimates of fitting polynomials or the results of the hypothesis tests: maximum likelihood fits are invariant under reparameterization and the estimated values of likelihood (and log likelihood) are unaffected by reparameterization (Mood et al., 1974, pp. 284ff).

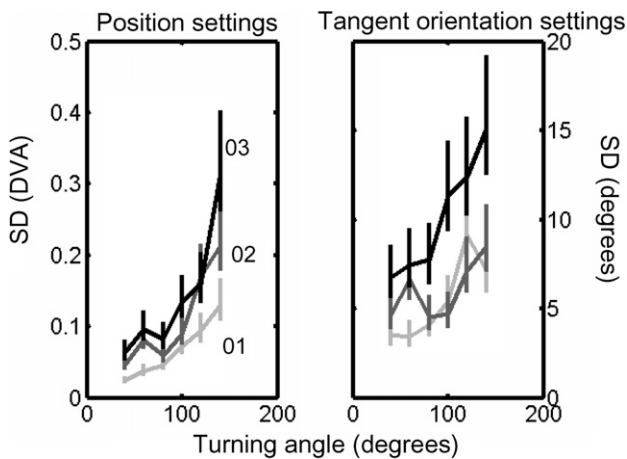


Fig. 5. Experiment 1a: Standard deviations. The standard deviations with 95% confidence intervals in position in terms of degrees of visual angle, DVA (left plot) and tangent orientation (right plot) settings for each of the three observers are shown.




Fig. 6. Experiment 1a: Hierarchy of hypothesis tests. (A) Hierarchy of nested hypothesis tests carried out for Experiment 1a. (B) Depiction of how the best-fitting polynomial was selected. See text for details.

Note that since Experiment 1a’s inducer pairs and the interpolants between them were all symmetric, the fits of interest were confined to polynomials of even degree. Carrying out the tests in order from left to right and bottom to top, the polynomial of the larger degree of the last rejecting test is the best-fitting polynomial to the observer’s data for the particular condition (see Fig. 6B). A summary of the results of the fitting procedure can be found in the upper half of Table 1 where the degrees of the best-fitting polynomials are listed, while Fig. 4 provides a graphical depiction of the best-fitting curves. Examining the results of the fitting procedure, we find that polynomials of higher degree are required to fit the data corresponding to inducer pairs with larger turning angles.

3.2.2.2. *Consistency.* Our next test of the influence of reliability investigates whether observers’ estimates of position and tangent orientation are consistent with any S^3 contour. This is similar in spirit to work by Koenderink and colleagues (see Koenderink, 1998; Koenderink, van Doorn, & Kappers, 1992) that studied shape from shading. They sought to test whether observers’ tangent orientation settings in a shape-from-shading task were consistent with any single, smooth surface. We will consider the internal consistency of observers’ interpolation settings, i.e., the extent to which an observer’s position and tangent orientation settings in a given condition are consistent with an S^3 contour.

We develop two measures of internal inconsistency of an observer’s settings: (i) a polynomial-based measure based on the degree of disagreement between observers’ orientation settings and the tangent orientations implied, via polynomial fit, by their position settings; and (ii) a non-parametric measure based on the degree of disagreement between observers’ position settings and the height differences implied by their orientation settings. In both cases, we obtain two different estimates of orientation or position settings; and we use a measure of inconsis-

tency that quantifies the degree of disagreement between the two.

3.2.2.3. *Polynomial-based measure of inconsistency.* Our first measure of inconsistency is based on the polynomial fits described in the previous section. The best-fitting polynomial is first computed to an observer’s position settings through the six interpolation windows in a given condition (i.e., for a particular inducer geometry). This polynomial predicts a tangent field, which is then compared against observers’ actual settings of tangent orientation through the six windows. The degree to which the observer’s orientation settings deviate from the tangents to the best-fitting polynomial reflects the extent of internal inconsistency of their settings.

Specifically, we have position (height) settings $h_{w,s}$ and orientation settings $\varphi_{w,s}$ through six window $w = 1, 2, \dots, 6$ each repeated $s = 1, \dots, 8$ times. Thus we have for each observer, 48 pairs $(h_{w,s}, \varphi_{w,s})$, $w = 1, \dots, 6$, $s = 1, \dots, 8$. First we reduce this data to 6 pairs by averaging across repetitions of setting:

$$h_w = \frac{1}{8} \sum_{s=1}^8 h_{w,s} \tag{3}$$

and

$$\varphi_w = \coprod_{s=1}^8 \varphi_{w,s} \tag{4}$$

where the symbol \coprod refers to the circular mean of angles (Mardia & Jupp, 2000). We find the best-fitting polynomial $P(w)$ to the position data h_w using maximum likelihood as described above (recall that the best-fitting polynomial is constrained to match the positions and orientations of the two inducing contours at their respective points of occlusion). The tangent orientations predicted by this polynomial at the six windows are given by:

$$\varphi'_w = \arctan P'(w), \quad w = 1, \dots, 6 \tag{5}$$

Table 1
Summary of Experiment 1a results for each observer

Observer	Turning angle					
	40	60	80	100	120	140
	Degree of best-fitting polynomial					
O1	2	4	4	6	6	6
O2	6	2	6	6	6	6
O3	2	4	4	4	4	6
	Polynomial-based inconsistency measure (Γ_P)					
O1	0.057	0.01	0.082	0.159	0.304	0.376
O2	0.531	0.491	0.734	0.109	0.425	0.64
O3	0.146	0.687	0.825	0.865	0.821	1.07
	Non-parametric inconsistency measure (Γ_{NP})					
O1	0.037	0.056	0.097	0.206	0.334	0.537
O2	0.385	0.396	0.363	0.198	0.149	0.218
O3	0.104	0.456	0.711	0.596	0.99	1.128

We wish to compare the actual orientation settings made by the observer $\varphi = (\varphi_1, \dots, \varphi_6)$ and the orientation settings implied by their location settings via the polynomial fit $\varphi' = (\varphi'_1, \dots, \varphi'_6)$. To do so, we first remove the mean from each vector,

$$\Delta\varphi = \varphi - \bar{\varphi} \tag{6}$$

where $\bar{\varphi} = \sum_{w=1}^6 \varphi_w / 6$. We define $\Delta\varphi'$ similarly. If we did not remove the means in this way, the resulting measure of inconsistency defined below would not be comparable across conditions and would change arbitrarily with changes in angular coordinate systems.

We define a measure of inconsistency patterned on the ratio of the squared residual error to the total variance in regression commonly referred to as “variance accounted for”. The analogue to the residual error is $\|\Delta\varphi - \Delta\varphi'\|^2$ where the double braces denote the vector norm. However, either $\|\Delta\varphi\|^2$ or $\|\Delta\varphi'\|^2$ could plausibly be used to normalize $\|\Delta\varphi - \Delta\varphi'\|^2$ and we use the average of the two, $(\|\Delta\varphi\|^2 + \|\Delta\varphi'\|^2)/2$. The resulting parametric inconsistency statistic based on the polynomial fits is

$$\Gamma_P = 2 \frac{\|\Delta\varphi - \Delta\varphi'\|^2}{\|\Delta\varphi\|^2 + \|\Delta\varphi'\|^2} \tag{7}$$

The factor of 2 in Eq. (4) serves no useful role and we omit it from our computation. In final form,

$$\Gamma_P = \frac{\|\Delta\varphi - \Delta\varphi'\|^2}{\|\Delta\varphi\|^2 + \|\Delta\varphi'\|^2} \tag{8}$$

We interpret small Γ_P values as evidence for a high level of internal consistency. Conversely, large values of Γ_P suggest inconsistency between an observer’s settings of position and tangent orientation—indicating that no S^3 interpolating contour could be consistent with the observer’s settings for that inducing-contour pair. We will use resampling methods (Efron & Tibshirani, 1993) to determine the distribution of Γ_P under the null hypothesis of consistency. This will allow us to test the hypothesis of consistency.

3.2.2.4. Non-parametric measure of inconsistency. Our second measure of consistency is based on deriving estimates of height differences from observers’ orientation settings at the six interpolation windows, and then comparing these with the height differences seen in the observers’ position settings. The degree to which the two estimates of height differences differ reflects the extent of internal inconsistency of observers’ settings.

Specifically, we have mean settings of position h_w and tangent orientation φ_w at the six windows, $w = 1, \dots, 6$. In addition, we also have virtual settings of position and tangent orientation corresponding to the two inducers at their respective points of occlusion. We denote these by h_0, h_7 and φ_0, φ_7 , respectively. We transform these settings into two sets of estimates of height differences on the seven intervals $(w - 1, w)$, $w = 1, \dots, 8$ —one based on the position set-

tings, the other based on the orientation settings. The height differences based on the position settings are simply:

$$\partial h_w = h_w - h_{w-1}, \quad w = 1, \dots, 7 \tag{9}$$

In computing the height differences implied by the orientation settings, we interpolate each adjacent pair of orientation settings to estimate the orientation setting at the center of each interval. We average pairs of angles from adjacent windows using the circular mean:

$$\bar{\varphi}_w = \prod_{v=w-1}^w \varphi_v, \quad w = 1, \dots, 7. \tag{10}$$

Then we compute the height difference on each interval, consistent with the interpolated angle,

$$\partial h'_w = \delta \tan \bar{\varphi}_w, \quad w = 1, \dots, 7 \tag{11}$$

As a result, we have two different sets of estimates of height differences, $\partial h = (\partial h_1, \dots, \partial h_7)$ and $\partial h' = (\partial h'_1, \dots, \partial h'_7)$ and we normalize both by subtracting their means. If we did not then the measure below would depend on the arbitrary choice of coordinate system. Let $\Delta h = \partial h - \sum_{w=1}^7 \partial h_w / 7$ and $\Delta h' = \partial h' - \sum_{w=1}^7 \partial h'_w / 7$. We quantify the degree of disagreement between the two using the following measure:

$$\Gamma_{NP} = \frac{\|\Delta h - \Delta h'\|^2}{\|\Delta h\|^2 + \|\Delta h'\|^2} \tag{12}$$

where we compare the vector magnitude of the differences to the average of the vector magnitude of the measures (omitting a factor of 2). Γ_{NP} corresponds to our second measure of inconsistency. As with Γ_P , we interpret small Γ_{NP} values as evidence for a high level of internal consistency, whereas large values indicate inconsistency between an observer’s settings of position and tangent orientation. As with the parametric inconsistency measure Γ_P , we will use resampling methods (Efron & Tibshirani, 1993) to determine the distribution of Γ_{NP} under the null hypothesis of consistency. This will allow us to test the hypothesis of consistency using the non-parametric measure.

There are other possible measures of inconsistency than the ones we have chosen. The chosen measures, Γ_P and Γ_{NP} , have the benefits that they are patterned on analogous measures in linear regression, familiar to many readers, and, as there are two of them, we can use each as a check on the other.

Table 1 shows the values of the two inconsistency measures, Γ_P and Γ_{NP} , for each observer’s settings under the six turning-angle conditions. We estimated the variability around the consistency value in each condition by classic bootstrap simulation with replacement (Efron & Tibshirani, 1993). The resulting estimates were used in computing confidence intervals ($\pm 1SD$; see Fig. 7). We used these SD values to test the null hypothesis that increasing turning angle does not affect the degree of inconsistency of observers’ settings. We formed a linear contrast of the six values of Γ_P (or Γ_{NP}) in order of increasing turning angle with weights




Fig. 7. Experiment 1a: Tests of consistency. For each condition and observer, we plotted the value of two inconsistency measures: polynomial-based (upper plot) and non-parametric (lower plot). The error bars represent ± 1 standard deviation, which was obtained through a bootstrapping procedure. See the text for details regarding these analyses.

$-3, -2, -1, 1, 2, 3$. We computed the $SD_{contrast}$ from the weights and the estimated SD 's of the six Γ_p were estimated by resampling. Under the null hypothesis of no trend the resulting contrast has expected value 0 and we can test the hypothesis by a two-tailed z-test (the SD s estimated by resampling do not need a correction for degrees of freedom).

The prediction of a reliable increase in inconsistency with increasing turning angle is generally not borne out (see Fig. 7). The polynomial-based measure Γ_P exhibited no significant increase with turning angle ($z_{\Gamma_P} = -0.7765, p = 0.219$ for a two-tailed test); whereas the non-parametric measure exhibited only a marginal effect ($z_{\Gamma_{NP}} = -2.0124, p = 0.022$). Inconsistency values are generally low across the range of turning angles tested, indicating a high degree of internal consistency: observers' settings are consistent with an S^3 interpolating contour in each turning-angle condition.

3.3. Discussion

In Experiment 1a we tested interpolation performance with a number of turning angles between indu-

cer pairs. The results indicate that there is no particular turning-angle cutoff that marks the boundary between visual interpolation and no visual interpolation—only a gradual, roughly linear increase in SD (i.e., decrease in precision) with increasing turning angle. Consequently, while we conclude that turning angle affects the precision of interpolation performance, we find that there is no evidence to support the claim of a discontinuity in performance as turning angles exceed 90° .

This conclusion was further supported by the analysis of internal consistency. Even for the most non-relatable turning angle, observers' settings of position and tangent orientation remained mutually consistent—indicating that visual completion was still occurring, contrary to the predictions of the reliability hypothesis. Indeed, this strong agreement between the fitted polynomials and the position and tangent orientation data validates these conclusions as something more than a direct consequence of the use of polynomials. We find little evidence for a decrease in consistency, graded or abrupt.

There is precedent in the literature that parabolas may play a special role in contour completion (e.g., Singh & Fulvio, 2005; Warren et al., 2002). This, in conjunction with the fact that all of the stimuli of Experiment 1a were symmetric and defined a unique parabola, encouraged the prediction that parabolas may provide acceptable fits to the data. As demonstrated above, this prediction was not borne out, but rather, observers' settings were consistent with curves that are substantially flatter than parabolic in most (15/18) cases.

We noted above that the inducer pairs in this experiment were co-circular at their respective points of occlusion, and the consistency of observer's performance is qualitatively consistent with a co-circularity hypothesis. While an increase in turning angle leads to lower precision, observers' settings nevertheless remain consistent with an S^3 contour.

4. Experiment 1b

The stimuli in Experiment 1a were all symmetric about the midline and co-circular. In Experiment 1b, we used asymmetric inducer pairs that again vary in their turning angle. The inducers in each pair had different orientations, thereby shifting the point of intersection of their linear extensions away from the horizontal midpoint of the occlusion region. A subset of the turning angles tested in Experiment 1a was included. If the stimuli of Experiment 1a were "special" because of their symmetric nature or co-circularity, one would predict a sharp deterioration in performance (i.e., precision, consistency) for the non-symmetric, non-relatable inducer pairs of Experiment 1b. Additionally, we would also expect worse performance with increasing deviation from symmetry.

4.1. Methods

4.1.1. Inducers

The turning angle between the inducer pairs took on one of three values: 60°, 80°, 100°. Within each pair, the inducers began with the same orientation and vertical position as if they were symmetric, and then each was rotated so that their intersection point shifted away from the horizontal midpoint of the occlusion region. The offset of the intersection point will be referred to as *skew*, denoted by γ . No relative vertical offsets were applied to the inducers (recall Fig. 3A right display).

For each of the three turning angles, two values of γ were used: $\gamma_1 = w/6$ and $\gamma_2 = w/3$ where w corresponds to the width of the rectangular occluder. Note that the $w/6$ intersection point is closer to the horizontal midpoint of the occluder than the $w/3$ intersection point because skew is measured relative to the midpoint.

4.1.2. Design

The design comprised three turning angles, two levels of skew, two sides of the occluder at which the reference inducer could be presented, and six window locations. Each session thus contained $3 \times 2 \times 2 \times 6$ or 72 trials, with each trial obtaining paired settings of position and tangent orientation of the line probe. Each observer performed adjustments in 4 such experimental sessions, preceded by a practice session.

4.2. Results and analysis

4.2.1. Analysis of precision

Fig. 8 shows the standard deviations for both setting types with 95% confidence intervals, for each observer, in each of the six conditions. In most cases, there is a general trend toward an increase in standard deviation with increasing turning angle. However, as in Experiment 1a, there is little indication of a change in performance at the postulated cutoff of 90°. Symmetry does not seem to be an important factor for the perceptual strength of visual interpolation, as the skew manipulation did not produce any systematic effects on standard deviation (i.e., the SDs were not significantly higher for the skewed inducer configurations than the symmetric configuration).

4.2.1.1. Analyses of shape and consistency. As in Experiment 1a, the shapes of the visually interpolated contours were first characterized by fitting polynomials to the position data only. The fitting procedure used was the same as in Experiment 1a, except that the polynomials were no longer constrained to have even degree. The best-fitting polynomials are shown superimposed on the data in Fig. 9, and their degrees are listed in the upper portion of Table 2. These fits exhibited an increase in the degree of polynomial fits with increasing deviation from symmetry in most (7/9) cases, and a small tendency toward increasing degree of polynomial fit with increasing turning angle.

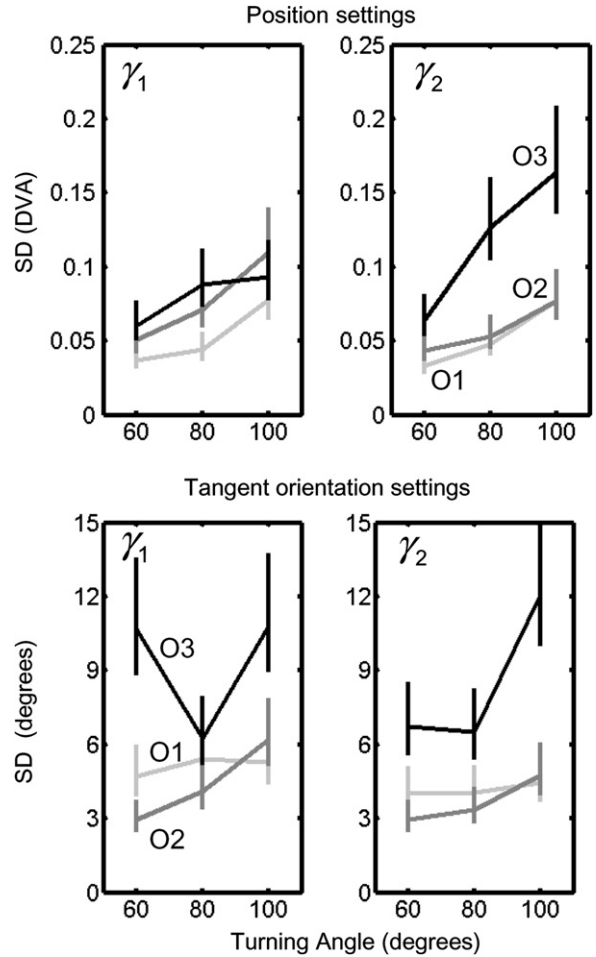


Fig. 8. Experiment 1b: Standard deviations in position and tangent orientation settings. The standard deviations with 95% confidence intervals in position in terms of degrees of visual angle, DVA (upper plots) and tangent orientation (lower plots) settings for each of the three observers are shown. The left column corresponds to the small skew (γ_1) condition, and the right column corresponds to the large skew condition (γ_2).

In Experiment 1a, observers’ settings of position and tangent orientation were found to exhibit a high degree of internal consistency in all conditions—including inducer pairs with non-relatable turning angles. The values of the two inconsistency measures, Γ_P and Γ_{NP} , for observers’ settings in Experiment 1b are shown in Table 2. As before, we estimated the variability around the inconsistency values in each condition by classic bootstrap simulation with replacement (Efron & Tibshirani, 1993). Confidence intervals ($\pm 1SD$) using these estimates are shown in Fig. 10. Statistical tests using these estimates revealed no significant effect of turning angle on the inconsistency of the settings for either level of skew. This was true for both the polynomial-based measure ($z_{\Gamma_P(\gamma_1)} = -1.6062, p = 0.054; z_{\Gamma_P(\gamma_2)} = -0.3116, p = 0.378$) and the non-parametric measure ($z_{\Gamma_{NP}(\gamma_1)} = -1.777, p = 0.039; z_{\Gamma_{NP}(\gamma_2)} = -0.0042, p = 0.50$). When testing the effect of skew on inconsistency, the polynomial measure showed a significant effect ($z_{\Gamma_P} = -2.0092, p = 0.022$) but the non-parametric measure did not ($z_{\Gamma_{NP}} = -0.6365, p = 0.262$).

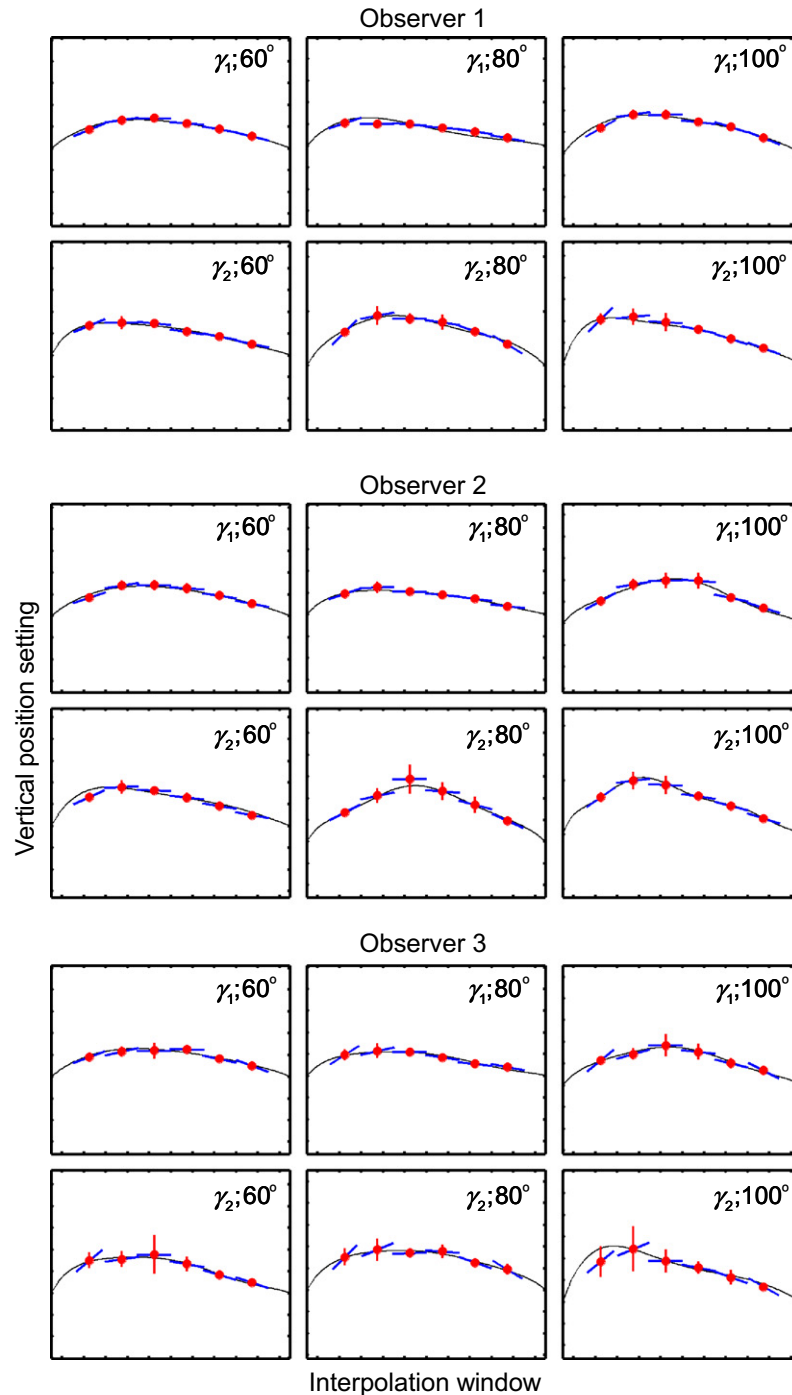


Fig. 9. Experiment 1b: Means and standard deviations in position and tangent orientation settings with the best-fitting polynomials. The format is identical to that of Fig. 4.

The manipulation of skew thus has but a weak effect of internal consistency—one that is not robust across the two measures of inconsistency.

4.3. Discussion

Experiment 1b comprised a second test of the role of turning angle in visual interpolation. The results with asymmetric inducer pairs were comparable to those of Experiment 1a; increasing turning angles led to lower pre-

cision in interpolation but there was no abrupt decrease at the 90° cutoff point. There was little or no effect of turning angle on internal consistency.

5. Experiment 2

We next focused on the second reliability criterion, which is the hypothesis that the linear extensions of the two inducers must intersect for interpolation to occur. As noted above, this criterion is equivalent to the exis-

Table 2
Summary of Experiment 1b results for each observer

Observer	Turning angle + skew					
	60 γ 1	80 γ 1	100 γ 1	60 γ 2	80 γ 2	100 γ 2
	Degree of best-fitting polynomial					
O1	4	4	4	6	7	6
O2	4	6	6	5	8	8
O3	4	6	6	6	6	5
	Polynomial-based inconsistency measure (Γ_P)					
O1	0.113	0.203	0.436	0.533	0.49	1.804
O2	0.251	0.78	0.678	0.757	0.625	0.735
O3	0.454	0.75	1.26	1.911	2.672	1.188
	Non-parametric inconsistency measure (Γ_{NP})					
O1	0.13	0.207	0.363	0.227	0.141	0.38
O2	0.22	0.554	0.607	0.262	0.261	0.333
O3	0.344	0.716	0.816	0.839	0.654	0.617




Fig. 10. Experiment 1b: Tests of consistency. The results of the polynomial-based (upper plots) and non-parametric (lower plots) consistency tests for each of the three observers. The left column corresponds to the small skew (γ_1) condition, and the right column corresponds to the large skew condition (γ_2).

tence of a smooth, non-inflecting contour that can interpolate between the two inducers. Since the relative vertical offset of a pair of inducers determines whether their linear extensions will intersect, we manipulated this variable in order to generate inducer pairs with different levels of

reliability. We are testing whether the failure of the linear extension intersection criterion, as measured by vertical offset, affects the precision and consistency of interpolation.

As in Experiment 1a, the inducers in a given pair had the same length and orientation. Three vertical offsets were used between the inducer pairs: one leading to symmetric relatable inducers (S), one to asymmetric but relatable inducers (R), and one to non-relatable inducers (NR), as defined below.

5.1. Methods

5.1.1. Inducers

We began with inducer pairs that had one of two relatable turning angles: 60°;90°. To each of these, we applied one of three possible vertical offsets: $\Delta = 0$ (symmetric, as in Experiments 1a & 1b), $\Delta = (2/3)w \tan \theta$ (relatable) $\Delta = (4/3)w \tan \theta$ (non-relatable; recall Eq. (1)).

5.1.2. Design

The design contained two turning angles, three vertical inducer offsets, two sides of the occluder on which the reference inducer could be presented, and six window locations. Each session thus contained $2 \times 3 \times 2 \times 6$ or 72 trials, with each trial obtaining paired settings of position and tangent orientation of the line probe. Each observer performed adjustments in 4 such experimental sessions, preceded by a practice session.

5.2. Results and analysis

5.2.1. Analysis of precision

Fig. 11 shows the standard deviations with 95% confidence intervals for the three observers as a function of vertical offset Δ , with the upper plots showing the SDs of the position settings, and the lower plots showing the SDs of the tangent orientation settings. The left column shows

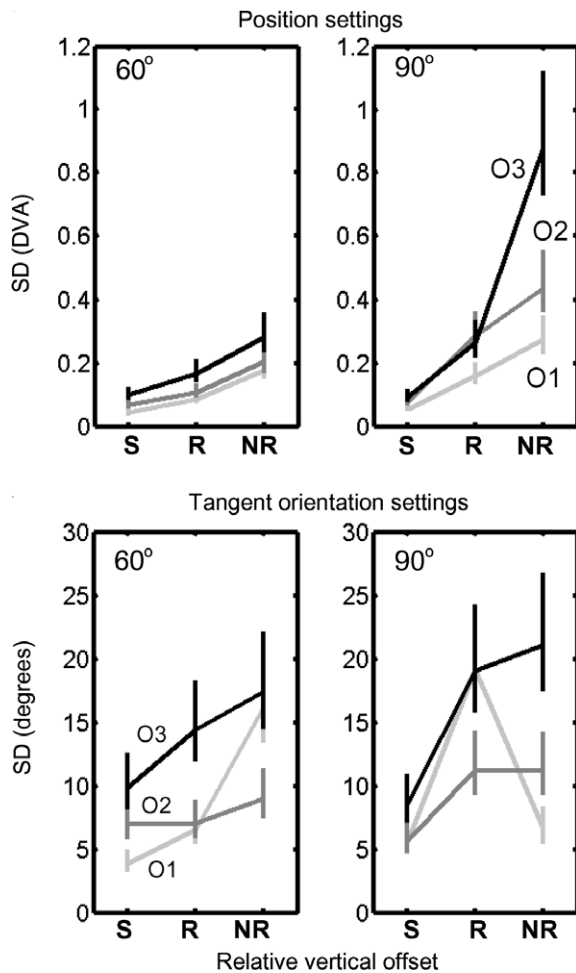


Fig. 11. Experiment 2: Standard deviations in position and tangent orientation settings. The format is identical to that of Fig. 8. Here, the left column corresponds to the 60° turning angle conditions, and the right column corresponds to the 90° turning angle conditions.

the SDs for the 60° turning angle conditions, and the right column shows the SDs for the 90° turning angle conditions. The results show that standard deviations increase with increasing Δ , as predicted (with one exception, O1 across the 90° turning angle conditions). There is also an increase in standard deviation with increasing turning angle (left vs. right plots) as seen in Experiments 1a and 1b, but observed mainly for the R and NR conditions.

5.2.2. Analysis of shape and consistency

As the vertical offset between a pair of inducers is increased, at some point inflecting curves become necessary to interpolate them, thereby creating unfavorable conditions for contour interpolation (Takeichi et al., 1995). The point at which this happens is precisely the point at which the linear extensions of the two inducers fail to meet (Singh & Hoffman, 1999). Thus, by both Takeichi et al's and Kellman & Shipley's theory, the non-relatable (NR) conditions in Experiment 2 should not be conducive for visual interpolation.

As in previous experiments, the shapes of the visually interpolated contours were characterized by fitting polynomials to the position data. The fitting procedure used was the same as in Experiment 1b (with no constraints on the degree of the polynomials). The results of the shape and consistency analyses for Experiment 2 are summarized in Table 3, and the best-fitting polynomials are shown superimposed on the data in Fig. 12. Confidence intervals (± 1 SD) around the inconsistency values (Γ_P and Γ_{NP}), based on classic bootstrapping stimulation with replacement, are shown in Fig. 13.

Statistical tests revealed a significant effect of relative vertical offset on the degree of inconsistency in observers' settings for both levels of turning angles. This was true for both the polynomial-based measure ($z_{\Gamma_P(60)} = -2.6504$, $p = 0.004$; $z_{\Gamma_P(90)} = -4.956$, $p < 0.001$) and the non-parametric measure ($z_{\Gamma_{NP}(60)} = -2.7969$, $p = 0.003$; $z_{\Gamma_{NP}(90)} = -2.8940$, $p = 0.002$). When testing the effect of turning angle on inconsistency, the parametric measure revealed a significant effect ($z_{\Gamma_{P(TA)}} = -3.3385$, $p < 0.001$), but the non-parametric measure did not ($z_{\Gamma_{NP(TA)}} = -0.3660$, $p = 0.357$).

5.3. Discussion

The predictions of relatability theory for the stimuli of Experiment 2 were supported, with observers' settings exhibiting an increase in internal inconsistency with increasing vertical offset. Our results provide direct evidence that failure of the inducer-extension intersection criterion of relatability has a detrimental effect on visual interpolation.

6. General discussion

Kellman and Shipley (1991) proposed two relatability criteria intended to characterize when visual interpolation of occluded contours would or would not occur. The first criterion was based on the turning angle between the inducers, the second on the intersection of their linear extensions (or lack thereof). The influence of the relatability criteria on shape interpolation has not been directly tested experimentally, in part because there is no obvious standard in the literature for what pattern in observers' interpolation settings counts as a failure to interpolate, and in part because they are stated as all-or-none criteria. We first reformulated both criteria as claims about smooth changes in interpolation performance controlled by the degree to which each criterion was violated (see also Singh & Hoffman, 1999).

Our study includes the first systematic tests of the influence of these relatability criteria on shape interpolation, based on three measures of successful interpolation that we developed. The first measure is a measure of imprecision based on the variability of settings across repeated trials which we take as a gauge of how precisely the visual system can localize a partly-occluded contour. The other

Table 3
Summary of Experiment 2 results for each observer

Observer	Relative vertical offset + turning angle					
	S60	R60	NR60	S90	R90	NR90
	Degree of best-fitting polynomial					
O1	2	3	4	6	6	4
O2	2	4	4	6	3	6
O3	4	4	5	6	4	4
	Polynomial-based inconsistency measure (Γ_P)					
O1	0.046	0.039	2.516	0.072	0.93	10.767
O2	0.078	0.137	0.813	0.146	0.898	4.427
O3	0.043	0.444	3.659	0.519	0.689	9.642
	Non-parametric inconsistency measure (Γ_{NP})					
O1	0.043	0.336	1.093	0.187	0.205	2.088
O2	0.105	0.2151	0.417	0.196	0.653	1.26
O3	0.42	0.682	4.535	0.448	0.363	4.235

two measures are indices of inconsistency across multiple settings. Observers are asked to judge the position and tangent orientation of an occluded contour at many points. We test whether each observer's position and orientation settings are consistent with any S^3 contour, or whether they are mutually inconsistent.

In three experiments, we measured observers' interpolation performance under a variety of inducer pair conditions that met or violated the reliability criteria. Experiments 1a and 1b examined the turning angle cutoff criterion, and Experiment 2, the inducer extension intersection criterion.

Based on the outcomes of these experiments, we conclude, first, that the turning angle between inducers affects the precision of interpolation. As the turning angle increased, the variability of observers' settings also increased, suggesting that the visual system is less precise in localizing an interpolating contour that undergoes greater turning (hence has higher curvature). This increase in setting variability (SD) was gradual, roughly linear, across the entire range of turning angles tested—both reliable and non-reliable. We found no abrupt, qualitative change in precision at 90° as predicted by Kellman & Shipley, or at any other angle. Moreover, neither measure of inconsistency was systematically affected by increased turning angle. The internal consistency of observers' position and orientation settings at multiple measurement locations remained high (i.e., we obtained low values of the inconsistency measures) for all turning angles tested—both reliable and non-reliable.

When we tested the second of Kellman & Shipley's reliability criteria, the inducer extension intersection criterion, we found both a decrease in precision (an increase in setting variability) and a decrease in internal consistency (an increase in both of our inconsistency measures) when large relative vertical offsets between inducers precluded intersection of their linear extensions.

How is it that large vertical offsets of the inducing edges lead to degradation in internal consistency of the visually interpolated contour, whereas large turning angles do

not? The answer may relate to Takeichi et al.'s (1995) claim that the strength of visual interpolation is inversely related to the number of inflections required by a smooth interpolating contour. Once the vertical offset between inducers exceeds the range that allows their extensions to meet, any smooth interpolating curve requires at least one inflection. This is not necessarily true with changes in inducer geometry involving turning angle only. Our results therefore support the hypothesis that failures of internal consistency occur when the interpolated contour requires a point of inflection.

Thus far, we have claimed that a failure of consistency implies that the observer is not interpolating an S^3 contour across the multiple measurements (i.e., position and tangent orientation settings through the 6 interpolation windows) for a given inducer geometry. Given the observers' performance in non-reliable (NR) conditions, we are disinclined to think that the observer is making settings at random, unable to make sense of the experimental instructions. Instead, we propose that the inconsistencies measured in the non-reliable offset conditions are due to the visual system completing multiple contours or contour fragments (e.g., individual extrapolants from the two inducers) that depend upon the location within the occluded region where the measurements are obtained, and that are influenced by the presence of the setting probe within the window. That is, failures of the extension intersection criterion lead to interpolation performance that is consistent with no S^3 contour across setting locations and setting types.

An examination of the data in the non-reliable (NR) offset conditions of Experiment 2 (Fig. 14), for example, reveals a tendency for observers' settings to be more variable in the central interpolation windows and less variable in the windows closest to either inducer. Thus, distance from the two inducers plays an important role in determining what is perceived at a given location. When the location is close to either inducer, the information from the closer inducer dominates, and is extrapolated to that location.

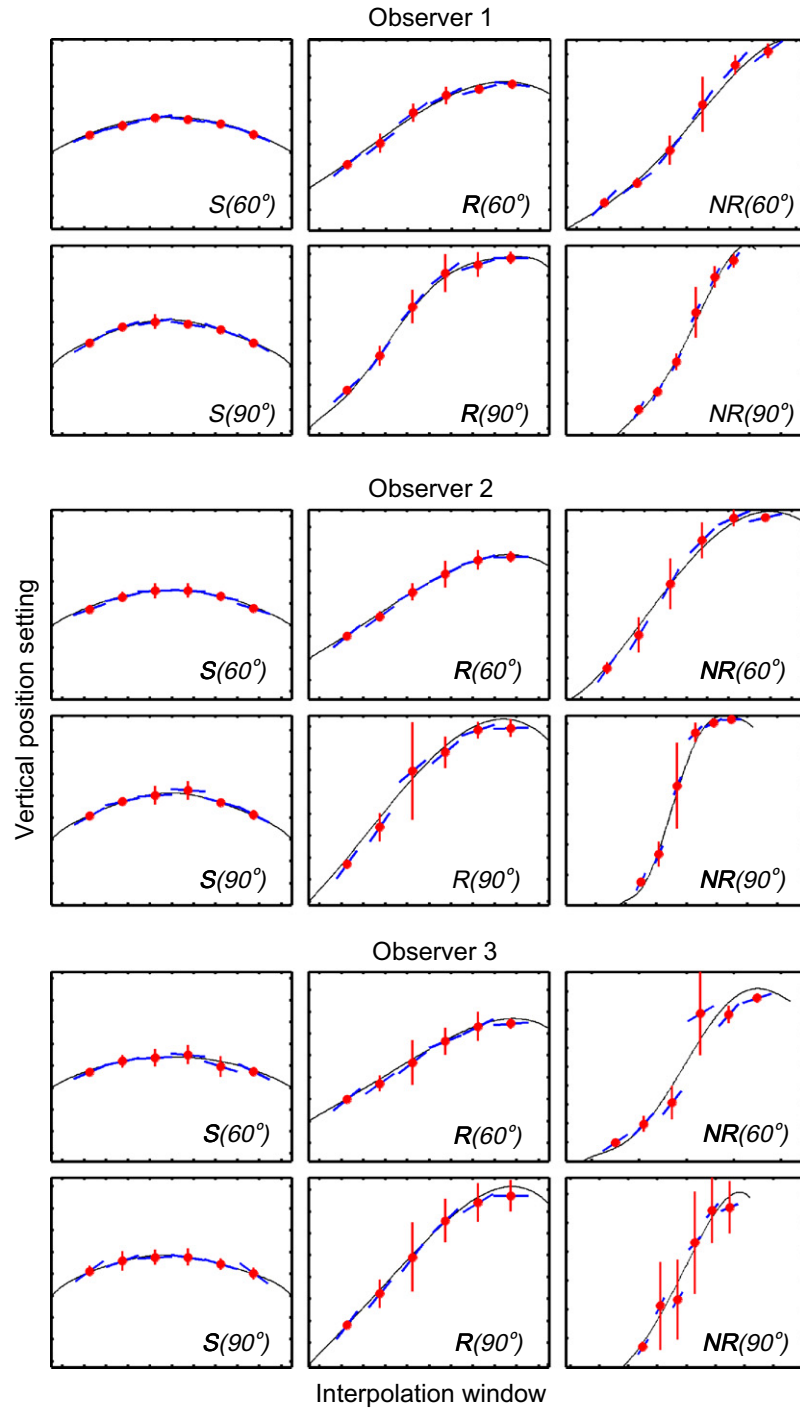


Fig. 12. Experiment 2: Means and standard deviations in position and tangent orientation settings with the best-fitting polynomials. The format is identical to that of Figs. 4 and 9.

This is not the case for the central interpolation windows, where the extrapolants of the two inducers define very different contours, and do not meet smoothly without the presence of an inflection. The large error bars on the mean settings made at these central locations in Fig. 14 could reflect mixtures between competing contours perceived on different setting trials.

Implicit in many discussions of contour completion is the hypothesis that the invisible portion of the contour is,

in some sense, explicitly represented in the visual system. We refer to this assumption as the “isomorphism hypothesis.” In experiments where observers make a single type of interpolation setting (e.g., position) on occluded contours, or make settings at a single location along an occluded contour, it is difficult to imagine what pattern of responses would be inconsistent with this hypothesis. However, in the experiments we report here, where observers make both orientation and position settings at multiple points along

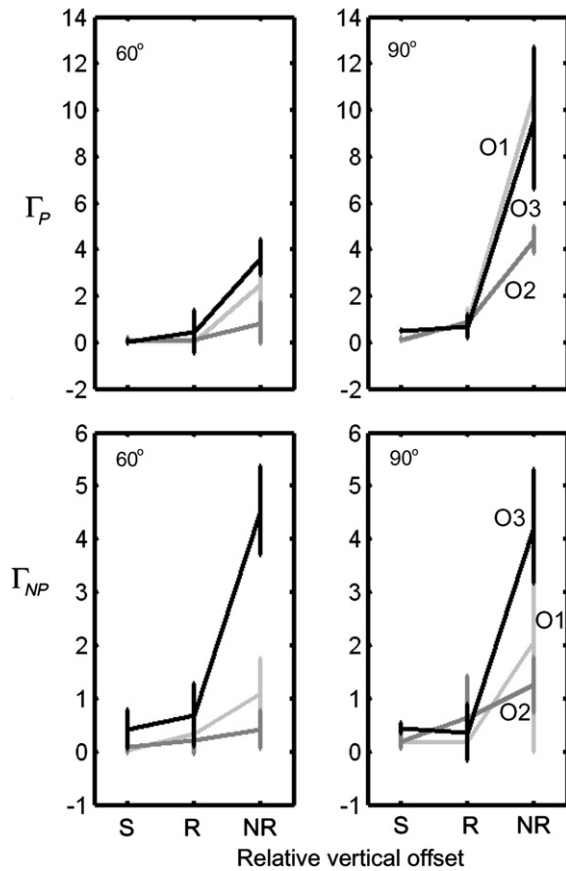


Fig. 13. Experiment 2: Tests of consistency. The format is identical to that of Fig. 10. Here, the left column corresponds to the 60° turning angle conditions, and the right column corresponds to the 90° turning angle conditions.

the occluded portion of the contour, it is possible for them to make settings that are consistent with no single, stable, smooth contour.

In those experimental conditions where the observer’s performance is markedly inconsistent, we must qualify the isomorphism hypothesis. We conjecture that, in such conditions, there is a representation of a complete contour

available on each trial, but it changes with type of judgment and location of judgment along the contour. As a thought experiment, we might imagine an observer viewing an occluded contour trying to decide where to make judgments along the length of the contour and whether to estimate its orientation or position, with the “isomorphic contour” changing with changes in intent.

We have sought to determine what geometric properties of the stimulus configuration affect the degree to which the continuation of an occluded contour is “good” or “bad”. Motivated by the notion of reliability and the association-field model, we investigated the role of two continuous properties—the turning angle between inducers and the intersection of their linear extensions. This allows us to replace a subjective criterion—good continuation—by two objective values. But the evident question now is: Why should these two geometric properties affect the precision and consistency of human visual interpolation as they do?

A plausible route to an explanation is to consider the importance of successful interpolation on an organism’s success in identifying partly-occluded objects, breaking camouflage, and manipulating such objects. Even as simple an action as grasping an opaque object between thumb and fingers involves completion of the backside of the object where the fingers will make contact without immediate visual guidance. Prior work has demonstrated the advantages of examining the statistics of natural images in understanding the geometry of contour integration (Elder & Goldberg, 2002; Geisler et al., 2001; Sigman et al., 2001). It is similarly likely that there is interesting structure in the statistics of partly occluded contours, and that it correlates well with human performance in completing partly occluded contours (Geisler & Perry, 2006). Specifically, it is plausible that variations in the precision of visual interpolation reflect actual uncertainty in occluded contours in our environment.

It is less clear how to relate failures of internal consistency to factors in the environment. A failure of consis-

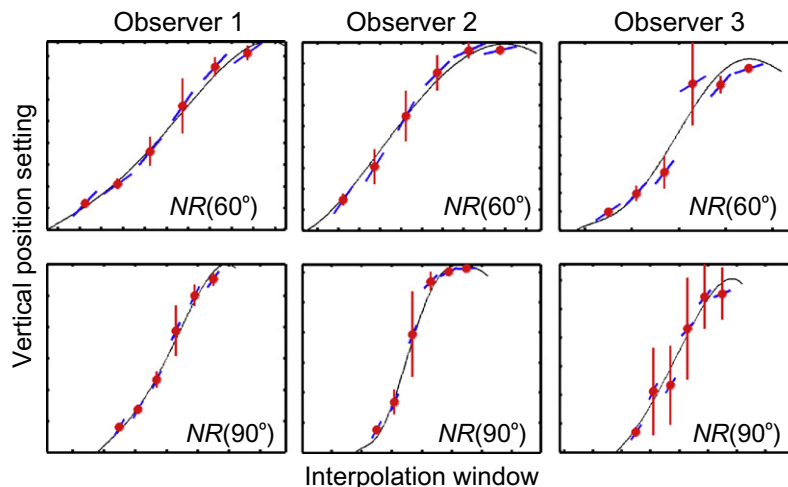


Fig. 14. Experiment 2: Data from the non-relatable, offset conditions. See text.

tency in our terms implies that a series of visual judgments correspond to no single state of the environment, to no ‘single, stable, smooth contour.’ We found, for example, that estimates of tangent orientation and of position in Experiment 2 were often consistent with no single contour. However, the judgments considered (position and tangent orientation) correspond to different visual tasks. In terms of Bayesian decision theory the organism, in carrying out a task, acts so as to minimize loss given particular choices of prior, likelihood and loss function (Knill & Richards, 1996; Maloney, 2002). It is plausible that performance in each task is optimized according to a separate loss function. In grasping for example, a failure to arrive at the proper grasp point can have different consequences than an error in orienting the finger pads to the surface. There is no obvious reason why judgments corresponding to tasks with different loss functions should be mutually consistent (Maloney, 2002, pp. 168ff). Maloney points out that the visual estimates of ideal Bayesian decision makers carrying out multiple tasks need not (indeed, should not) correspond to any consistent “pictorial” representation of the environment.

In summary, we conjecture that the precision of visual interpolation of occluded contours is grounded in the statistical environment with imprecision reflecting the objective uncertainty concerning the occluded parts of contours. The reliability criteria of Kellman & Shipley capture properties of the environment that affect the precision with which occluded contours can be localized or their tangent orientations judged. In considering the internal consistency of visual judgments, and the conditions under which visual judgments are mutually consistent or inconsistent, we broach a potentially deeper question. It is not obvious that the visual judgments of human observers correspond to any single “pictorial” representation of the environment (the “isomorphism hypothesis” discussed above)—or that they should. Further research on this point would likely result in a better understanding of how visual information including visual uncertainty is represented in the brain and how it is used in performing perceptual tasks and planning action.

Acknowledgments

This research was funded in part by a Graduate Research Fellowship from the National Science Foundation (J.M.F.), Grant BCS-0216944 from the National Science Foundation (M.S.), and Grant EY08266 from the National Institute of Health (L.T.M.). We thank Wilson S. Geisler for extensive and useful comments on a previous draft.

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.
- Ben-Shahar, O., & Zucker, S. (2004). Geometrical computations explain projection patterns of long-range horizontal connections in visual cortex. *Neural Computation*, *16*, 445–476.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.
- Davis, G., & Driver, J. (1998). Kanizsa subjective figures can act as occluding surfaces at parallel stages of visual search. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(1), 169–184.
- Efron, B., & Tibshirani, R. (1993). *An introduction to the bootstrap*. New York: Chapman-Hall.
- 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.
- 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 and Psychophysics*, *63*(7), 1171–1182.
- 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. M., & 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.
- Geisler, W. S., & Perry, J. S. (2006). Efficiency of contour grouping across occlusions in natural images Abstract. *Journal of Vision*, *6*(6), 336, 336a. doi:10.1167/6.6.336http://journalofvision.org/6/6/336/ .
- Grossberg, S., & Mingolla, E. (1985). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Reviews*, *92*, 173–211.
- Guttman, S. E., & Kellman, P. J. (2004). Contour interpolation revealed by a dot localization paradigm. *Vision Research*, *44*, 1799–1815.
- Guttman, S. E., Sekuler, A. B., & Kellman, P. J. (2003). Temporal variations in visual completion: A reflection of spatial limits?. *Journal of Experimental Psychology: Human Perception and Performance* *29*(6), 1211–1227.
- Hon, A. K., Maloney, L. T., & Landy, M. S. (1997). The influence function for visual interpolation. *Proceedings of the SPIE: Human Vision and Electronic Imaging II*, *3016*, 409–419.
- Kellman, P. J., Guttman, S. E., & Wickens, T. D. (2001). Geometric and neural models of object perception. In T. F. Shipley & P. J. Kellman (Eds.), *From fragments to objects: Segmentation and grouping in vision* (pp. 183–246). Amsterdam, The Netherlands: Elsevier Science.
- Kellman, P. J., & Shipley, T. F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology*, *23*, 141–221.
- Knill, D. C., & Richards, W. (Eds.). (1996). *Perception as Bayesian inference*. Cambridge, UK: Cambridge University Press.
- Koenderink, J. J. (1998). Pictorial relief. *Philosophical Transactions of the Royal Society of London, A*, *356*, 1071–1086.
- Koenderink, J. J., van Doorn, A. J., & Kappers, A. M. (1992). Surface perception in pictures. *Perception and Psychophysics*, *52*, 487–496.
- Liu, Z., Jacobs, D., & Basri, R. (1999). The role of convexity in perceptual completion: Beyond good continuation. *Vision Research*, *39*, 4244–4257.
- Maloney, L. T. (2002). Statistical decision theory and biological vision. In D. Heyer & R. Mausfeld (Eds.), *Perception and the physical world: Psychological and philosophical issues in perception* (pp. 145–189). New York: Wiley.
- Mardia, K. V., & Jupp, P. E. (2000). *Directional statistics*. New York: Wiley.
- Michotte, A., Thinès, G., & Crabbé, G. (1964). *Les compléments amodaux des structures perceptives*. Louvain, Belgique: Publications Universitaires de Louvain.

- Mood, A. M., Graybill, F. A., & Boes, D. C. (1974). *Introduction to the theory of statistics* (3rd ed.). New York: McGraw-Hill.
- Parent, P., & Zucker, S. W. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *11*(8), 823–839.
- Pelli, D. G. (1997). The Video toolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Pizlo, Z., Salach-Golyska, M., & Rosenfeld, A. (1997). Curve detection in a noisy image. *Vision Research*, *37*(9), 1217–1241.
- Rauschenberger, R., & Yantis, S. (2001). Masking unveils pre-amodal completion representation in amodal completion. *Nature*, *410*, 369–372.
- Ringach, D. L., & Shapley, R. (1996). Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision Research*, *36*(19), 3037–3050.
- Sekuler, A. B., & Palmer, S. E. (1992). Perception of partly-occluded objects: A microgenetic analysis. *Journal of Experimental Psychology: General*, *121*(1), 95–111.
- Shipley, T. F., & Kellman, P. J. (1992). Perception of partly occluded objects and illusory figures: Evidence for an identity hypothesis. *Journal of Experimental Psychology: Human Perception and Performance*, *10*, 106–120.
- Sigman, M., Guillermo, A. C., 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), 1035–1040.
- Singh, M. (2004). Modal and amodal completion generate different shapes. *Psychological Science*, *15*, 454–459.
- Singh, M., & Hoffman, D. D. (1999). Completing visual contours: The relationship between relatability and minimizing inflections. *Perception and Psychophysics*, *61*, 943–951.
- Singh, M., & Fulvio, J. M. (2005). Visual extrapolation of contour geometry. *Proceedings of the National Academy of Sciences USA*, *102*, 939–944.
- Singh, M., & Fulvio, J. M. (2007). Bayesian contour extrapolation: Geometric determinants of good continuation. *Vision Research*, *47*, 783–798.
- Takeichi, H., Nakazawa, H., Murakami, I., & Shimojo, S. (1995). The theory of the curvature-constraint line for amodal completion. *Perception*, *24*, 373–389.
- Takeichi, H. (1995). The effect of curvature on visual interpolation. *Perception*, *24*, 1011–1020.
- Warren, P. E., Maloney, L. T., & Landy, M. S. (2002). Interpolating sampled contours in 3D: Analyses of variability and bias. *Vision Research*, *42*, 2431–2446.
- Warren, P. E., Maloney, L. T., & Landy, M. S. (2004). Interpolating sampled contours in 3D: Perturbation analyses. *Vision Research*, *44*, 815–832.
- Wertheimer, M. (1923). Untersuchungen zur Lehre von der Gestalt, II. *Psychologische Forschung*, *4*, 301–350.
- Yin, C., Kellman, P. J., & Shipley, T. F. (2000). Surface integration influences depth discrimination. *Vision Research*, *40*, 1969–1978.