

Consistency of location and gradient judgments of visually-interpolated contours

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

Manish Singh
Psychology & Cognitive Science
Rutgers, New Brunswick
Piscataway, NJ 08854
manish@rucss.rutgers.edu

Laurence T. Maloney
Psychology & Neural Science
New York University
New York, NY 10003
ltm1@nyu.edu

Abstract

We report two experiments assessing how observers interpolate partly-occluded contours induced by pairs of line segments that disappeared behind an occluder. On each trial observers iteratively adjusted the location and orientation of a short line probe that could be moved vertically within the region of occlusion. They were instructed to set the line probe to be tangent to the occluded contour. The line probe could appear at one of six horizontal offsets and the inducer pairs on some trials were relatable (the inducers could be joined by a smooth curve without a point of inflection) and on other trials non-relatable. We interpreted the settings as estimates of the location and gradient (slope) of the contour in the region of occlusion. We tested whether the resulting visual estimates of location and gradient were consistent with any single smooth contour. When inducers were relatable, estimates of location and gradient were mutually consistent for all observers and could be modeled as polynomials of 5th or lower degree. When the inducers were non-relatable, the consistency of location and gradient settings deteriorated, indicating that there can be no single smooth curve that accounts for observers' judgments. We discuss the implication of these results for models of human visual interpolation.

1. Introduction

Perceptual organization allows us to perceive objects and people, rather than an incoherent mosaic of luminance and color [20]. An important aspect of perceptual organization is the completion of object boundaries for which there is only partial evidence in the visual field, based on their "good continuation" [21].

Visual completion of contours consists of three conceptually separate tasks—*detection* of contour information in the scene, correct *grouping/segmentation* of the elements of each contour, and contour *interpolation/extrapolation* [19, 16]. It is likely that all three tasks are carried out concurrently in human perceptual organization with grouping, for example,

affecting detection, etc. However, it is conceptually convenient to divide the process into these three 'tasks'.

The first two tasks have been studied extensively (e.g. [2, 4, 5, 7]) and recent experimental work addresses the third [16, 19, 18, 14, 6]. A number of models have been proposed to interpolate the shape of missing portions of contours, based on geometric [3, 17], variational [12, 9], and stochastic [22] approaches.

The issue we address here grows out of this previous work. It is best introduced by imagining that we ask an observer to make many different kinds of judgments concerning an occluded contour. We could ask an observer to identify points within the region of occlusion that the contour passes through. We could also ask for estimates of the local orientation or curvature of the contour at each point it passes through. We could ask observers to locate qualitative properties of the contour such as the locations of extrema of curvature. Once we have gathered this information, we could compare it to model predictions for any of the models mentioned above. But there is a prior question that we should first address: how do we know that there is any single visual contour that is consistent with all of the observer's judgments?

Koenderink and colleagues [10, 11] addressed an analogous issue in the perception of smooth surfaces represented in pictures. An observer, asked to set gradient probes to be tangent to surfaces at many points, will do so. Koenderink and colleagues tested whether the settings for each observer were consistent with any smooth surface in three-dimensional space. They integrated the depth changes implicit in gradient settings along closed contours in the picture plane. The result of such integration on any smooth surface must be 0 (no net change in depth). If any integral on a closed contour consistently deviated from 0, then there would be no smooth surface consistent with the observer's data. Koenderink and colleagues could not reject the hypothesis that observers perceived a single, stable smooth surface in each picture they viewed and were making estimates of its gradient.

Here we examine whether observer's estimates of location and gradient of a partly-occluded contour in the region of occlusion are consistent with any smooth

contour. We describe the observer's tasks next and then describe how we tested consistency.

2. Experiment

The experimental strategy we employ is based on one used recently to measure the shape of visually-extrapolated contours [16]. This involves (i) obtaining *paired* settings of interpolation position and orientation (rather than position alone), and (ii) obtaining settings at a number of locations along the length of a partly-occluded contour, in order to derive a relatively detailed map of the *extended* shape of the visually-interpolated contour.

2.1 Observers

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




Figure 1: Schematic of the basic stimulus display used in the experiment.

2.2 Apparatus

The basic stimulus display is shown in Figure 1. It was composed of two oriented line segments of the same fixed length, referred to as the *inducers*. The inducers were placed at the left and right edges of an occlusion region marked by a gray rectangle. Its width is denoted w . From condition to condition, we varied the angle $\theta \in \{30^\circ, 45^\circ\}$ of each inducer from the horizontal, and the vertical offset between the inducers at their respective points of occlusion, denoted Δ .

On a given trial, a small vertical slit 0.14 degrees of visual angle in width appeared at one of six locations within the occluder, which we refer to as an *interpolation window*. All six are shown in Figure 1, although, on any

given trial, only one would appear. The vertical midlines of the interpolation windows were 0.42 degrees of visual angle apart, and the leftmost and rightmost windows were 0.49 degrees of visual angle 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 location and orientation were to be adjusted by the observer. All stimuli were presented in the center of the screen on a black background.

Inducers: The inducers had a visible length of 3.34 degrees of visual angle and had one of two turning angles between them: 60° or 90° , which corresponded to individual orientations of 30° or 45° relative to the horizontal, respectively. The inducers were white and 0.028 degrees of visual angle thick, and were anti-aliased at the resolution of one-fourth of a pixel. The inducers pointed upwards 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 location was used as the reference for vertical placement of the opposite inducer (see explanation in the next paragraph). On any given trial, the reference inducer could appear either on the left or on the right.

Reliability: If the two inducers in Figure 1 can be joined by a smooth curve without a point of inflection, with a total turning angle of no more than 90° , we refer to them as *reliable*, otherwise they are *non-reliable* [8, 15]. Inducer reliability thus provides a prediction for when grouping of scene fragments should occur. Moreover, such grouping has direct implications for interpolation—strongly-grouped inducers should elicit more precise and mutually-consistent settings of interpolation.

Given that the two inducers make the same angle with the horizontal, it can be shown that the inducers are reliable if and only if

$$|\Delta| < w \tan \theta \quad (1)$$

as illustrated in Figure 2. One of three possible vertical offsets was applied between the two inducers on any trial:

$$\Delta = 0 \quad (\text{symmetric, reliable (S)}), \quad \Delta = \frac{2}{3} w \tan \theta$$

$$(\text{reliable (R)}) \quad \Delta = \frac{4}{3} w \tan \theta \quad (\text{non-reliable (NR)}).$$




Figure 2: Demonstration of reliable inducers.

Occluders: The rectangular occluder had a height of 9.66 degrees of visual angle and a width of 3.06 degrees of visual angle. On each trial, one of six interpolation windows of width equal to 0.14 degrees of visual angle appeared within the rectangle.

Adjustable probe: A straight-line probe was visible through the window, which had the same color, thickness, and anti-aliasing as the inducers. The probe was initially presented with a horizontal orientation, at the vertical midpoint of the interpolation window. In the height-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 height, denoted h , was constrained to lie between the uppermost and lowermost 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.

Display: The stimulus displays were presented on a high-resolution 22-in. monitor (*Lacie Blue*), in 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. All stimuli were displayed using the PsychToolbox extensions for MATLAB [1], [13].

2.3 Design and Procedure

Each observer ran 4 experimental sessions, preceded by a practice session. Within a session, each pair of inducers (thus, each combination of turning angle and relative inducer offset) was presented twice for each of the six window locations (with the reference inducer presented once on the left and once on the right). As indicated above, the design contained two turning angles, three vertical inducer offsets, two sides of the occluder at

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.

Two of the experimental sessions presented the inducing contours pointing upwards while the other two presented them pointing downwards. Their order was counterbalanced. On any given trial, observers first adjusted the height 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 fashion between height and orientation settings in order to optimize the percept of a smooth partly-occluded contour between the two inducing contours. They pressed the trackball button to move on to the next trial when they were satisfied with the combination of positional and orientation settings.

3 Results

For each trial we recorded the observer's positional ('height') and orientation settings at each of the six windows, which we treated as estimates of the location and gradient, respectively, of the partly-occluded contour in the region of occlusion. Following preliminary analysis, we were able to combine conditions that differed only in left-right reflection. We will describe the stimuli and settings as if all stimuli were transformed to common coordinates. For each of the remaining six conditions (2 values of θ , 3 values of Δ) we computed the mean height settings $\bar{h}_1, \dots, \bar{h}_6$, and mean orientation settings $\bar{\phi}_1, \dots, \bar{\phi}_6$ for each of the interpolation windows. Figure 3 shows the interpolation data for one of the observers, in each of the six inducer-pair conditions. In what follows, we report various analyses performed on these data.

3.1 Analysis of Consistency

We determined the extent to which the positional and orientation settings in a given inducer condition are consistent with a single smooth curve. First, we computed $\bar{\delta}_i^h = \bar{h}_i - \bar{h}_{i-1}$ the estimated differences between interpolation curve heights defined by the mean positional settings at successive window locations. There were seven differences computed ($i = 1, 2, \dots, 7$), with \bar{h}_0 and \bar{h}_7 defined as the heights of the left and right inducers, respectively, at their respective points of occlusion.




Figure 3: Interpolation data—mean position and orientation settings—for one observer for each inducer pair condition.

Next, we computed $\tilde{\delta}_i^\phi, i=1, \dots, 7$, the estimated difference between curve heights at successive windows based on mean orientation settings,

$$\tilde{\delta}_i^\phi = s \left[\left(\tan(\bar{\phi}_i) + \tan(\bar{\phi}_{i-1}) \right) / 2 \right] \quad (2)$$

where s is the horizontal separation between the midlines of two successive windows, and $\bar{\phi}_0 = \theta$ and $\bar{\phi}_7 = -\theta$ are the orientations of the left and right inducers, respectively.

High internal consistency implies that the height and orientation settings conform to a common smooth curve. Thus, $\tilde{\delta}_i^h$ and $\tilde{\delta}_i^\phi$ should be roughly equal. We defined a measure of consistency as $IC = 1 - D^2 / \hat{\sigma}^2$, where values near 1 imply high internal consistency and values near 0 imply low internal consistency. D^2 was computed as

$$D^2 = \sum_{i=1}^7 \left(\tilde{\delta}_i^\phi - \tilde{\delta}_i^h \right)^2 / 7 \quad (3)$$

and $\hat{\sigma}$ is the standard deviation of $\tilde{\delta}_i^\phi$.




Figure 4: Results of the internal consistency analysis for one typical observer. Internal consistency is high for reliable inducers, but declines sharply for non-reliable inducers.

Typical results for one observer are plotted in Figure 4. The internal consistency is evidently high for reliable inducers, and declines sharply for non-reliable inducers. We will quantify these observations below. These results suggest that the reliability or non-reliability of the inducers affects the perceptual ‘strength’ of partly-occluded contours. Visual-interpolation performance declines drastically when this criterion is not satisfied.

The turning angle criterion of reliability requires that the total turning angle between inducers not exceed 90° . In order to test this, interpolation measurements were

obtained on an additional set of stimuli. This set contained inducer pairs that had $\Delta = 0$ and took on one of six different turning angles: 40° , 60° , 80° , 100° , 120° , 140° . By the turning-angle criterion, the inducer pairs with the last three turning angles should yield interpolants with dramatically lower internal consistency, reflecting the decreased perceptual strength of interpolation.

The results of the internal consistency analysis for these conditions for all three observers are shown in Figure 5. Bearing in mind that smaller values of *IC* imply lower internal consistency, the results are inconsistent with the turning-angle criterion. One observer exhibited no systematic change in internal consistency with turning angle, whereas the other two observers showed an overall decline in internal consistency with increasing turning angle. In neither case, was there a marked drop in internal consistency at the postulated 90° cutoff—rather a uniformly gradual decrease was observed throughout the range of turning angles tested.




Figure 5: Results of the internal consistency analysis for all three observers for the extended set of stimuli used to further test the turning angle criterion of reliability. There is no systematic effect of the postulated 90° inducer turning angle cutoff of the reliability hypothesis.

3.3 Precision

The precision of the settings was analyzed as another measure of the strength of the percept of the interpolating curve. The reliability hypothesis of Kellman & Shipley [8] leads to the prediction that those inducer pairs that satisfy the two criteria for reliability should group strongly, thereby yielding substantially more precise settings of the interpolating contour. Figure 6 depicts the standard deviations for one typical observer as a function of Δ . In the left plot are shown the standard deviations for the positional settings for the two turning angle

conditions. In the right plot are shown the standard deviations for the orientation settings for the two inducer turning angle conditions. For both setting types, standard deviation increases with (i) increasing Δ ; and (ii) increasing turning angle (for $\Delta \neq 0$).




Figure 6: Standard deviation in position (left graph) and orientation (right graph) settings for one typical observer. Standard deviation increases with increasing Δ and inducer turning angle.

Figure 7 shows the standard deviations plotted as a function of turning angle between inducers for the second set of stimuli with the larger range of turning angles (but with Δ fixed at 0) for all three observers. Consistent with the above findings, standard deviation increases with turning angle for both position and orientation settings. There does not appear, however, to be a marked effect at the 90° reliability cutoff, but rather an overall increase with turning angle. Larger turning angles led to weakening of the perceptual ‘vividness’ of interpolated contours.

4 Shape Modeling

In order to characterize the human visual spline, we fit polynomials to the data. We forced all polynomials to match the slope and location of the inducers at the endpoints adjacent to the points of occlusion (four constraints). An n 'th degree polynomial has $n+1$ free parameters and consequently, we could satisfy the constraints imposed on the polynomial when $n \geq 3$. However, when $\Delta = 0$ (the symmetric case) we could also fit a quadratic polynomial ($n=2$) that satisfied the inducer constraints. When $n > 3$, the imposed constraints do not determine the resulting polynomial but specify a family of polynomials. We indexed the resulting families of polynomial curves for $n > 3$ by adding $n-3$ additional height constraints at *control points* spaced equally across the occlusion regions. These control points function as free parameters that we could adjust in matching

polynomials of a specified degree to observers' height settings.




Figure 7: Standard deviation in position (left graph) and orientation (right graph) for all three observers for the extended set of stimuli used to further test the turning angle criterion of reliability. As with internal consistency, there is no systematic effect of the postulated 90° inducer turning angle cutoff of the reliability hypothesis.

For $n \leq 3$, we calculated the sum of the squared differences (SS) between the (unique) polynomial and the observer's height settings. For $n > 3$, we numerically varied control points to select the single polynomial in the family that minimized the SS.

Since only the positional settings were taken into account in the polynomial fits, the predicted orientation settings based on those fits serve as a test of how well these fits also characterize the orientation settings. These analyses were carried out by a simple comparison of the predicted orientations settings (the slope of the fitted polynomial at each height across the occluded region) and obtained orientation settings. These comparisons showed that the orientation settings for the symmetric inducer interpolants are highly-consistent with those predicted by the best-fitting quintic polynomials for those conditions. As the inducers become increasingly offset, however, the agreement between the observers' orientation settings and those predicted by the best-fitting polynomials deteriorates, with the settings made in the non-relatable inducer conditions being markedly different from those predicted by the best-fit curves. For non-relatable inducers, there can be no smooth curve that matches observers' estimates of both location and gradient.

In the relatable case, polynomials of lower than degree 5 can provide good fits to both height and orientation data in certain special cases. However, in general, a quintic polynomial is required to match both orientation and height settings. Figure 8 demonstrates failures of fits occurring with polynomials of order less than 5. This result is especially surprising in the case of symmetric inducers ($\Delta = 0$), in that parabolas (i.e., quadratic polynomials) fail to capture the shape of interpolating contours even in this case. The interpolation settings in

this case (recall Figure 3) follow a contour that is substantially "flattened" with respect to the best-fitting quadratic. As a result, at least a fourth-degree polynomial is needed to model the shape of the interpolating contour for symmetric, relatable inducers.

Figure 9 shows the polynomial fits to the relatable inducers (i.e. the S and R conditions). The first row depicts the quintic polynomial fits to one typical observer's height settings for each of the four relatable conditions, while the second row shows the comparison between predicted orientation settings and the observer's orientation settings. The S inducers are symmetric about a vertical axis and if the observers' settings are also close to symmetric, then the quintic will typically be very close to a quartic (the coefficient of the term x^5 will be close to 0) as is the case. The observer's orientation settings for the symmetric conditions are highly-consistent with those predicted by the best-fitting quintic. For offset, relatable inducers, polynomials of lower degree do not provide fits that capture the observers' height and orientation settings. Even with the quintic polynomial fits, the agreement between predicted and obtained orientation settings deteriorates somewhat. This is in agreement with the *internal-consistency (IC)* results discussed above (Fig 4).

Figure 10 shows the polynomial fits to the non-relatable inducers. As in Figure 9, the first row shows the quintic polynomial fit, and the second row plots the observer's orientation settings vs. predicted orientation settings from the polynomial fits. It is clear that the polynomial fits to the positional settings do not capture observers' orientation settings. Again, this trend is in agreement with the corresponding *IC* values.

5. Discussion

We asked observers to make paired settings of position and orientation of partly-occluded contours at six locations across the region of occlusion. The contours were induced by two line segments on either side of the occluding surface. When these inducers were relatable, observers' settings of position and orientation were consistent with one another. Moreover, in the relatable case, polynomials of 5th degree were sufficient to match perceived location and gradient settings. When the inducers were non-relatable, we found that observers' height and orientation settings were not mutually consistent. We conclude that that there is no single smooth curve consistent with human settings with the non-relatable inducers we employed.




Figure 8: Examples of fit failures of polynomials with degree less than 5. In the top picture, the attempted fit is quadratic. In the center picture, the attempted fit is cubic, and in the bottom picture, the attempted fit is quartic.

Our results confirm that the distinction between relatable and non-relatable inducers [8] is important in understanding how the visual system completes contours (or whether it can complete a given contour reliably). We note that, whereas the criterion that the connecting contour between two inducers must be smooth, without any points of inflection, proved critical for reliable interpolation in our study, the criterion that the total turning angle between the inducers not exceed 90° did not play such a role in our findings.

Our results also constrain possible models of human visual contour completion. Any candidate model for “the human visual spline” must, first of all, be able to closely approximate quintic polynomials for relatable contours. Second, any such model should include a plausible

explanation for the evident inconsistency in location and gradient settings in the non-relatable case.

References

- [1] D.H. Brainard. The Psychophysics Toolbox. *Spatial Vision*, 10: 433-436, 1997. 3
- [2] J.H. Elder and R.M. Goldberg. Ecological statistics of Gestalt laws for the perceptual organization of contours. *Journal of Vision*, 2(4):324-353, 2002. 1
- [3] C. Fantoni and W. Gerbino. Contour interpolation by vector-field combination. *Journal of Vision*, 3(4):281-303, 2003. 1
- [4] J. Feldman. Curvilinearity, covariance, and regularity in perceptual groups. *Vision Research*, 37: 2835-2848, 1997. 1
- [5] D.J. Field, A. Hayes, and R.F. Hess. Contour integration by the human visual system: Evidence for a local “association field.” *Vision Research*, 33:173-193, 1993. 1
- [6] J. M. Fulvio and M. Singh. Surface geometry influences the shape of illusory contours. *Acta Psychologica*, in press. 1
- [7] W.S. Geisler, J.S. Perry, B.J. Super, and D.P. Gallogly. Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41(6):711-724, 2001. 1
- [8] P.J. Kellman, and T.F. Shipley. A theory of visual interpolation in object perception. *Cognitive Psychology*, 23:141-221, 1991. 2,4,5
- [9] B. B. Kimia, I. Frankel, and A. Popescu. Euler spiral for shape completion. *International Journal of Computer Vision*, 54(1/2):157-180, 2003. 1
- [10] J. J. Koenderink. Pictorial relief. *Philosophical Transactions of the Royal Society of London, A*, 356: 1071-1086, 1998. 1
- [11] J. J. Koenderink, A. J. van Doorn, and A. M. Kappers. Surface perception in pictures. *Perception and Psychophysics*, 52: 487-496, 1992. 1
- [12] D. Mumford. *Elastica and computer vision*. In C. L. Bajaj, editor, *Algebraic geometry and its applications*, pages 491-506. Springer-Verlag, New York, 1994. 1
- [13] D.G. Pelli. The Video toolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10:437-442, 1997. 3
- [14] M. Singh. Modal and amodal completion generate different shapes. *Psychological Science*, 15:454-459, 2004. 1
- [15] M. Singh and D.D. Hoffman. Completing visual contours: The relationship between relatability and minimizing inflections. *Perception and Psychophysics*, 61:943-951, 1999. 2
- [16] M. Singh and J.M. Fulvio. Visual extrapolation of contour geometry. *Proceedings of the National Academy of Sciences USA*, 102:939-944, 2005. 1,2
- [17] S. Ullman. Filling-in the gaps: The shape of subjective contours and a model for their generation. *Biological Cybernetics*, 25:1-6, 1976. 1




Figure 9: Quintic polynomial fits to the height settings of reliable inducers (first row) and orientation setting comparisons between those predicted by the fits and those set by the observer (second row) for one typical observer. The observer's actual orientation settings are highly accurate with respect to those predicted, however, this agreement begins to deteriorate for fits to the offset, reliable conditions. Note the agreement of these trends with the values of the *IC* measure shown in parentheses.

- [18] P.A. Warren, L.T. Maloney, and M.S. Landy. Interpolating sampled contours in 3D: Analyses of variability and bias. *Vision Research*, 42:2431-2446, 2002. 1
- [19] P.A. Warren, L.T. Maloney, and M.S. Landy. Interpolating sampled contours in 3D: perturbation analyses. *Vision Research*, 44:815-832, 2004. 1
- [20] M. Wertheimer. Untersuchungen zur Lehre von der Gestalt II. *Psychologische Forschung*, 4:301-350. 1
- [21] M. Wertheimer. Experimentelle Studien über das Sehen von Bewegung. *Zeitschrift für Psychologie*, 60: 321-378, 1912. 1
- [22] L. R. Williams and D. W. Jacobs. Stochastic completion fields: a neural model of illusory contour shape and salience. *Neural Computation*, 9:837-859, 1997. 1




Figure 10: Quintic polynomial fits to the location settings of non-reliable inducers (first row) and orientation settings comparisons between those predicted by the fits and those set by the observer (second row) for one typical observer. There is a strong lack of agreement between the predicted orientation settings from the curve fit and those made by the observer, indicating that there is no smooth curve consistent with this observer's settings. Note the agreement of these trends with the values of the *IC* measure shown in parentheses.