

AN EXPERIMENTAL CRITERION FOR CONSISTENCY IN INTERPOLATION OF PARTLY-OCCLUDED CONTOURS

JACQUELINE M. FULVIO¹ MANISH SINGH³ LAURENCE T. MALONEY^{1,2}

¹Department of Psychology
New York University

²Center for Neural Science
New York University

³Department of Psychology and Center for Cognitive Science
Rutgers University

In press, *Journal of Vision*, December 10, 2008

Correspondance:

Jacqueline M. Fulvio
New York University
Department of Psychology
6 Washington Place Rm. 275
New York, NY 10003

Email: jmf384@nyu.edu

Keywords: contour interpolation, consistency, isomorphism hypothesis, Poggendorff Illusion

Running head: Consistency in contour interpolation

ABSTRACT

We develop an experimental measure of consistency of interpolation of partly-occluded contours based entirely on observers' interpolation performance. We first describe the measure, which is based on a two-probe task that compares an observer's interpolation settings at a particular location with vs. without the observer's own setting presented at a nearby location. We then report two experiments aimed at investigating the behavior of the measure. The first compares the proposed measure to the predictions of contour completion models. The second investigates its performance in the Poggendorff and related configurations. We find that consistency covaries with reliability and cocircularity (both interpreted in graded terms) and, sensibly, yields a low measure for interpolation in the Poggendorff configuration. We conclude that the proposed measure of consistency operationalizes an important aspect of what is meant by the "strength" of a partly-occluded contour.

[139 words]

INTRODUCTION

The objects we perceive in everyday, cluttered scenes are often occluded in part by other objects. As a consequence, only a fraction of their boundaries is imaged onto the retinas. The visual system generates representations of contours and surfaces that go beyond fragmentary image data, and this information is part of the visual representations that underlie perception and guide action. Upon seeing a display such as in Figure 1A, typical observers do not suppose that they are looking at two dark-colored surfaces with an intervening light-colored surface. Rather they perceive a single dark-colored surface that is partly-occluded by, and continues underneath, the light-colored surface similar to the depiction in Figure 1B. This form of visual completion is referred to as “amodal” because, despite the vivid percept object unity, observers do not actually “see” a contour (or local contrast) in the occluded region (Michotte, Thinès & Crabbé, 1964/1991; Kanizsa, 1979). Amodal completion is generally contrasted against “modal” completion, such as in illusory contours, where observers do perceive a contour in image regions that contain no contrast.

Figure 1 about here

There is a considerable literature demonstrating that such amodally completed representations have measurable consequences for visual perception and attention (e.g., Sekuler & Palmer, 1992; Ringach & Shapley, 1996; Davis & Driver, 1998; Liu, Jacobs, & Basri, 1999; Rauschenberger & Yantis, 2001; Anderson, Singh, & Fleming, 2002; Guttman, Sekuler, & Kellman, 2003). A separate and prior question, though, is whether the visual system can complete a particular contour in a particular scene. Kellman & Shipley (1991), for example, proposed relatability criteria intended to predict when interpolation of partly-occluded contours would or would not occur. An evident difficulty with testing the relatability hypothesis is that explicit empirical criteria are

required for deciding whether the visual system has “successfully” interpolated the missing portions of partly-occluded contours.

Our goal in this paper is to propose and test an experimental criterion for “successful” interpolation of a partly-occluded contour based on the self-consistency of observers’ judgments about the contour.¹ We wish to complement existing experimental criteria based on observer ratings of “strength”, “goodness”, or “salience” of contours. It is not obvious what the experimenter means by “strength”, etc, or how these ratings relate to the observer’s performance in interpolating contours. In such experiments, the experimenter measures not what observers can do but instead uses a measure that is at best related to what observers believe they can do. (In giving a high completion rating, for instance, an observer is in effect expressing his/her confidence that, if asked to interpolate the missing contour, s/he could readily do so.) Although performance-based measures have previously been proposed, they typically rely on observers’ performance on some other task that is believed to rely on visual completion—for example, discriminating between two qualitative shape types (Ringach & Shapley, 1996), depth discrimination (Liu, Jacobs & Basri, 1999), or visual search (Davis & Driver, 1998; Rauschenberger & Yantis, 2001). However, Fulvio, Singh, and Maloney (2008) proposed two measures – setting precision and setting consistency – that were based on what observers actually do when they interpolate position and orientation settings along partly-occluded contours (see below). The criterion proposed here will also be based on observers’ interpolation performance, but will emphasize setting consistency.

Previous work on shape interpolation has for the most part focused on obtaining positional measurements at one particular location within the occluded region (e.g. Takeichi, 1995; Hon, Maloney, & Landy, 1997; Fantoni & Gerbino, 2003; Stanley &

¹ As will become clear in what follows, by “successful” interpolation we do not mean that the visually-interpolated contour matches the geometry of the actual contour that is occluded; but rather that the visual system generates a globally consistent representation of *some* smooth interpolated contour.

Rubin, 2003; Guttman & Kellman, 2004; Warren, Maloney, & Landy, 2002, 2004; Fulvio & Singh, 2006; but see Singh & Fulvio, 2005; Fulvio, Singh, & Maloney, 2008). It is difficult to see, however, what pattern of settings at a single location would support or refute a consistent representation. Thus, to test whether the visual system generates the representation of an internally-consistent interpolating contour, we consider multiple settings at different locations along the contour. With such settings, one can test whether the two (or more) types of observer settings at the different locations are in fact mutually consistent. We emphasize that nothing about our experiment requires that observers “see” contours in the occluded region; indeed, this follows from the very definition of amodal completion. They are being asked to interpolate the parts of a contour that are occluded.

Recently, Fulvio, Singh, & Maloney (2008) measured both the precision of contour interpolation and its “self-consistency”. Their approach to measuring consistency was inspired by previous work on shape from shading by Koenderink, van Doorn, & Kappers (1996). In the context of surface perception, Koenderink et al. (1996) had observers make repeated gradient settings at multiple locations on shaded surfaces. They then tested whether the observers’ judgments were consistent with any single, smooth surface. They did not reject the hypothesis of consistency, suggesting that observers’ percepts of surfaces are self-consistent and stable.

Fulvio et al. (2008) had observers make both orientation and position judgments at multiple locations along a partly-occluded contour (the stimuli were similar to those used in the first experiment reported here). They then tested whether the judgments were consistent with any single contour connecting the inducers. They proposed two different measures of consistency: a polynomial-based measure and a non-parametric measure. In computing the polynomial-based measure, for instance, they first fitted a polynomial to the positional settings alone, and then measured the extent to which the

tangent orientations derived from this polynomial predicted observers' actual settings of orientation. For co-circular and "relatable" inducer pairs, the settings exhibited a high degree of consistency. But for "non-relatable" inducer pairs requiring an inflecting interpolating contour, the degree of consistency was low, suggesting that the visual system does not generate a self-consistent representation of an interpolating contour under these conditions. Their proposed measures of consistency were in good agreement with cocircularity and Kellman & Shipley's (1991) relatability criterion that the linear extensions of the two inducing edges must intersect, if these criteria are interpreted in graded terms (see Fulvio et al., 2008). As such, the further the inducing contours are from cocircularity, the greater will be the asymmetry of the point of intersection of the linear extensions with respect to the points of occlusion, the weaker would be the predicted completion strength.

The key idea in the above approach to measuring consistency is to ask the observer to make many estimates of the location and orientation in the occluded portion of a partly-occluded contour and then to test whether the estimates, taken together, are consistent with any single, stable, smooth contour across the occluded region. The weaknesses of the method are the large number of settings needed and (for the polynomial version) the arbitrary introduction of a polynomial approximation.

In the current work, we describe a simple two-probe method that can be used to assess consistency. In Figure 2, we illustrate the method by means of a thought experiment (that departs from the methodology of the actual experiments described in detail below). In Figure 2A, we show a partly-occluded contour defined by two line segments on either side of a grey occluded region. The observer is asked to interpolate the partly-occluded contour in the region where it is occluded. We could ask the observer to tell us the vertical height of the contour for any horizontal displacement within the occluded region. The observer's settings at different horizontal displacements

are represented by a curve of *first settings* (red curve). We emphasize that the solid red line through the occluded region represents the locus of first settings, where the observer would set the interpolating point at each horizontal offset. As noted above, we do not assume that the observer in any sense “sees” this contour and, for reasons that will become immediately obvious, we do not assume that the curve of first settings¹ corresponds to the observer’s estimate of the full contour.

Figure 2 about here

Next, imagine that the experimenter changes the display in Figure 2A by adding a single point (probe) at some location along the curve defined by the observer’s first settings (Figure 2B). The observer is told that this point is part of the partly-occluded contour and, as it is his own setting, he is unlikely to disagree. We can now ask the observer to add a second point at any other horizontal displacement and thereby obtain a curve of *second settings*, conditional on the presence of the probe. If the curve of second settings falls on top of the curve of first settings, then the addition of the probe has not affected interpolation. If, however, the curve of second settings is different from the curve of first settings (as suggested by the blue dashed contour in Figure 2B) then it is evident that the probe itself has altered interpolation performance.

The two-probe method we develop and the new consistency criterion we propose rely on the following linking hypothesis: If observers represent a single, stable, smooth partly-occluded contour for a particular inducer pair configuration, their interpolation settings should be unaffected by the presence of their own setting at a nearby location. As a result, any systematic difference between the curve of first settings and the curve of second settings would provide evidence for a lack of internal consistency in the observer’s settings, and suggest that the observer fails to represent a single, stable,

¹ By the “first settings” we mean the settings obtained in this first part of the experiment.

smooth interpolated contour. The method and criterion will be tested in two experiments reported below. The experiments deviate in detail from the thought experiment just discussed, most notably in that we will ask observers to interpolate both location and orientation (tangent).

Based on the results of these experiments, we propose that consistency operationalizes (Bridgman, 1927) an important aspect of what is intended by the term “strength” when applied to partly-occluded contours. Moreover, it is a measure that is based solely on observers’ interpolation performance; it does not require that the observer explicitly rate “strength”, “goodness”, or “salience” of partly-occluded contours, nor does it rely indirectly on performance in some other task that is believed to depend in some way on interpolation.

The first experiment will use stimuli that are predicted by models of contour completion to produce stable or unstable representations—namely, those that can be interpolated smoothly only with an inflecting curve,—and have been demonstrated to do so by Fulvio et al. (2008). The second experiment will employ a configuration that is unlikely to produce a stable representation—the configuration that produces the Poggendorff illusion. We emphasize that although the contours and the illusion themselves are interesting in their own right, our primary focus in this article is to determine whether our consistency criterion and the two-probe method are suitable for the investigation of the stability with which interpolated contours are represented by the visual system.

Despite the limited experimental insight into the stability and consistency of representations of visually-completed contours and surfaces, there have been several attempts to characterize the nature of representations in general, held by the visual system. One hypothesis that may have bearing on the current study is the *isomorphism hypothesis*. The isomorphism hypothesis is, first of all, the claim that the visual

representation is effectively a three-dimensional model of the immediate environment that includes estimates of properties of objects in the scene including their positions, orientations, and material properties:

... objects are represented in a way that preserves the information most essential for survival—information about the inherent properties of objects, and about the organism's spatial relations to them (Shepard, 1981, p. 291)

But the isomorphism model is more than just a list of estimates of positions and properties. The organism uses the information in the model to make judgments about objects and to plan movement: to decide the best way to grasp the edge of a desk, to decide how best to move the desk through a doorway, or to decide whether the disorganized stack of papers on the desk will survive the move. Shepard used the term “second-order isomorphism” (Shepard & Chipman, 1971) to emphasize that the isomorphism model of the environment is a computational tool that permits the organism to anticipate the consequences of its own planned actions and other changes in the environment.

In the context of partial occlusion, this would entail constructing a self-consistent representation of the hidden portions of contours and surfaces. Thus, if the isomorphism hypothesis is accurate, observers should accept their own initial settings as part of the perceived interpolating-contour when making subsequent settings. If they do not, then the resulting inconsistency would call into question the isomorphism hypothesis. The latter finding would not be unprecedented as several previous studies have revealed inconsistencies and dissociations in the processing of multiple attributes of stimuli (e.g., Wenderoth, 1983; Gillam & Chambers, 1985; Smeets et al, 2002), thereby calling the isomorphism hypothesis into question. We also note that the inconsistencies found in the non-relatable conditions of Fulvio et al (2008) provide additional evidence against the isomorphism hypothesis. As described above, the two-

probe method developed below will only require that observers be consistent with themselves.

EXPERIMENT 1

Experiment 1 employed the “two-probe interpolation” method. As noted above, if observers represent a single, stable, smooth partly-occluded contour for a particular inducer pair configuration, their settings through an interpolation window should not be affected by the presence of their own setting in a nearby window. (One may, however, expect an increase in the reliability of the settings, given the additional information available in the two-windows task.)

Methods

Observers.

Eight observers at New York University completed the experiment. None were aware of the purpose. All had normal or corrected-to-normal vision.

Stimuli.

Examples of the stimulus displays are shown in Figure 3A,B. The basic components of the displays comprised two oriented line segments referred to as *inducers*. The inducers were white, 6.05 degrees of visual angle (DVA) long, and 5.8 arcmin thick, with anti-aliasing at the resolution of one-fourth of a pixel. They had equal orientations of 30° but with opposite signs. We denote orientation by θ and the width of the occluder by w (Figure 3C).

Figure 3 about here

The relative vertical offset between each pair of inducers at their respective points of occlusion, denoted Δ , was manipulated (see Figure 3C for a schematic depiction). Three levels of Δ were used: $\Delta = 0$ (“symmetric”, non-offset inducers);

$\Delta = \frac{2}{3} w \tan \theta$ (“small offset” inducers); $\Delta = \frac{4}{3} w \tan \theta$ (“large offset” inducers). We note

that by one of Kellman & Shipley’s (1991) *reliability* criteria, the large-offset inducers are non-relatable—their linear extensions do not intersect. Equivalently, the inducers in this case can be smoothly interpolated only with an inflecting curve (one that changes its sign of curvature); see Takeichi, Nakazawa, Murakami, & Shimojo, (1995); Singh & Hoffman, (1999). There is prior evidence that such inducer configurations are problematic for human visual interpolation (Kellman & Shipley, 1991; Takeichi et al., 1995; Fulvio et al., 2008).

The inducers pointed upward in half of the experimental trials, 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 (Δ) 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 the occlusion region, shown as a gray rectangle with rounded corners. It had width w equal to 6.05 DVA and height equal to 16.82 DVA. Narrow vertical slits of width 0.24 DVA, appeared within the occluder at one of four horizontal locations, referred to as an *interpolation window*. (See below for specific details concerning the presentation of the interpolation windows in each experimental part.) The vertical midlines of the windows were 0.87 DVA apart, and the leftmost and rightmost windows were 1.72 DVA from the closest vertical edge of the occluder. None of the windows appeared at the horizontal midpoint of the occluder.

Part I: Single Interpolation Window

In Part I, each trial contained one interpolation window through which a white, straight-line probe was visible (see Figure 3A). The probe’s 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.

Part II: Two windows

In Part II, the displays contained an interpolation window with the adjustable probe as in the Part I trials, as well as a second window that contained a line segment whose position and orientation were determined by the mean settings of that observer through that window in Part I. We reduced the possible pairs of windows that could be displayed to those we believed were most informative. This led to pairings of windows 1 & 3 and windows 2 & 4 (i.e. the mean settings appeared in window 3 and the adjustable probe in window 1 in 25% of the trials and vice versa, and likewise for the pairing of windows 2 & 4; see Figure 3B).

Software and apparatus.

The stimuli and experiments were programmed in MATLAB using the Psychophysical Toolbox extensions (Brainard, 1997; Pelli 1997). The computer used in the experimental apparatus was a Sony GDM-FW 900 workstation with a 24-inch

monitor with a display area of 48.3 cm x 30.5 cm and a resolution of 1600 by 1200 pixels at a vertical refresh rate of 75 Hz. The color quality was 32 bit and the graphics processor was a Quadro4 380x61. The system processor of the computer was an Intel Pentium 4 with SSE2. The stimuli were presented to the observers in the center of the screen upon a black (0.3 cd/m^2) background from a distance of approximately 127 cm from the computer screen.

Procedure.

As described above, on each trial observers viewed displays containing a rectangular surface whose vertical edges abut two linear inducers. The inducers could have one of three vertical offsets relative to each other. An interpolation window was opened within the rectangular occluder through which an adjustable line probe was visible. In Part II, there was also a second window through which a line segment was visible, whose position and orientation were determined by the observer's own mean settings from Part I. The task of the observers in both experimental parts was to adjust the position, h , and orientation, ϕ , of the interpolation probe until the combination of settings optimized the percept of a smoothly continuing partly-occluded contour defined by the two inducing contours. Observers first adjusted the position of the line probe vertically within the window, using a mouse. Pressing the space bar then allowed them to toggle between adjusting the position and orientation of the line probe. Observers toggled back and forth in this manner between height and tangent orientation settings. They were instructed to optimize the percept of a smooth partly-occluded contour defined by the two linear inducing contours. They clicked the mouse button when they were satisfied with the combination of height and tangent orientation settings.

Design.

The design for both experimental parts contained three inducer offsets and four window locations (4 possible single windows in Part I; 4 possible window pairs in Part II). The reference inducer could appear either on the left or the right side of the occluder, and the inducers could be oriented either upwards or downwards. Each session thus contained $3 \times 4 \times 2 \times 2$ or 48 trials, times three repetitions of each combination for a total of 144 total trials per experimental part. (In each experimental part, there were thus 48 trials for each inducer-geometry condition—12 per window.) As indicated, paired settings of position and tangent orientation of the line probe were obtained on each trial. Each observer performed adjustments in two experimental parts, one with a single interpolation window and one with two interpolation windows as described above, each preceded by a practice session to ensure understanding of the task and to allow the observer to become accustomed to the controls.

Results and analysis.

On each single-window trial (Part I), observers adjusted the position and orientation of the adjustable line probe so that it optimized the percept of a smoothly continuing contour between the two inducers. In the two-window trials (Part II), observers performed the same task with their own mean settings from one of the other windows in the single-window trials displayed in the stimulus. Left-right reflection and upwards-downwards inducer presentation did not influence observers' settings so we combined conditions and transformed them all into a common coordinate system. For each of the three values of inducer offset Δ , we computed the mean position settings h_1, \dots, h_6 , and mean tangent orientation settings ϕ_1, \dots, ϕ_6 , through each of the interpolation windows. Below we describe the results and analyses for each individual inducer-offset condition.

Symmetric inducers.

The symmetric inducers have the same relative height along the edges of the occluder ($\Delta = 0$), and are therefore symmetric about the vertical midline of the occluder. These inducers are also co-circular (Parent & Zucker, 1989): the two inducers are tangent to a common circle at their respective points of occlusion. The means and standard deviations of the settings for the symmetric inducers are shown in Figure 4: Part I settings are shown in red and Part II settings are shown in blue. To remind the reader how these settings were obtained in the context of the experiment, we have isolated one of the observer's pairs of Part I and Part II settings in Figure 4 (O2, window 2) and depicted them in terms of the experimental paradigm in Figure 5A. For this particular pair that we have isolated, the observer made the depicted mean settings in windows 2 & 4 in Part I (upper panel). Subsequently, in Part II, when the observer's own mean settings were displayed in window 4 (depicted in red in the lower display of Figure 5A), the observer made settings through window 2 whose mean is depicted in green. Superimposing the window 2 mean settings from Part I in black, we can see by comparison that the observer made almost identical settings with and without the probe in window 4. By contrast, in Figure 5B, we illustrate one observer's inconsistent setting for the large-offset condition (O5, window 3; see the "Large-offset inducers" section for more details).

Figure 4 about here

Figure 5 about here

Upon inspection of Figure 4, we see that this pattern is typical as the means generally do not appear to differ between Part I and Part II for symmetrical inducers. Indeed, *t*-tests on the differences between mean settings from Part I to Part II (8

observers x 4 windows = 32 tests for each setting type) were consistent with this as only 2/32 tests for changes in position were significant at the $\alpha = 0.05$ level, and only 1/32 test for changes in orientation was significant.

Since the settings evoked by symmetric inducers are consistent across time and space, the mean settings displayed in Part II should serve as an added cue to the percept. This should lead to lower setting variability with no change in the mean point of estimation (Landy, Maloney, Johnston, & Young, 1995; Ernst & Banks, 2002; Oruç, Maloney, & Landy, 2003; Boyaci, Doerschner, & Maloney, 2006; Fulvio, Singh, & Maloney, 2006). F-tests (8 observers x 4 windows = 32 tests for each setting type) on the ratio of the variances across the two experimental parts were performed. We found that 22/32 tests were significant for increases in setting precision (i.e., decreases in setting variability) for position settings and 12/32 tests were significant for increases in setting precision for orientation settings. Thus the added probe improves contour localization, but does not provide the same benefits for processing of the higher-order attributes (i.e. orientation) of the interpolating-contour. Nevertheless, only 4/32 tests were significant for decreases in orientation setting precision, suggesting that the added probe was rarely at odds with the perceived orientation. Although this and any other improvements in setting precision could be due to practice, we remind the reader that no feedback is provided to the observers that would motivate a change in strategy or performance as they make their settings. Additionally, as will be seen below, setting precision does not improve for the Part II settings in all conditions. Finally, the inducers were randomly perturbed along the edges of the rectangular occluder from trial to trial and their left-right and up/down presentation was random and counterbalanced. Thus, it is unlikely that observers are simply learning or remembering where to position the probe over time.

Large-offset inducers.

In the large-offset condition, the two inducers have been shifted sufficiently along the edges of the occluder so as to require an inflecting-contour for smooth interpolation. These inducers are thus “non-relatable”. The means and standard deviations of the settings for the large-offset inducers are shown in Figure 6: As in Figure 4, Part I settings are shown in red and Part II settings are shown in blue. As with the symmetric inducers, we have isolated one of the observer’s pairs of Part I and Part II settings in Figure 6 and depicted them in terms of the experimental paradigm in Figure 5B. For this particular pair that we have isolated, the observer made the depicted mean settings in windows 3 & 1 in Part I (upper panel). Subsequently, in Part II, when the observer’s own mean settings were displayed in window 1 (depicted in red in the lower display of the figure), the observer made settings through window 3 whose mean is depicted in green. Superimposing the window 3 mean settings from Part I in black, we can see by comparison that the observer’s Part II settings deviate strongly from the original settings for that window in Part I.

Figure 6 about here

As is evident in Figure 6, this pattern is typical in the large-offset condition: Unlike the symmetric-inducers condition, here there is visible deviation between observers’ settings in Part I and Part II. *t*-tests on the differences between mean settings from Part I and Part II yielded 24/32 tests that were significantly different at the $\alpha = 0.05$ level for position, and 8/32 tests for changes that were significantly different for orientation. Thus, interpolation position is strongly influenced by the presence of the observer’s own mean setting in a nearby window. (Although orientation settings are less affected than positional ones, they are nevertheless affected more so than for symmetric inducers.)

As noted above, a reliable difference between observers' settings between Part I and Part II indicates that there is no single, stable, smooth contour that is consistent with the observers' settings. There is, however, a definite, consistent pattern in the deviations observed between Part II and Part I. The settings in windows 1 & 2 are shifted downwards in the presence of the line probe in windows 3 & 4 respectively, and the settings in windows 3 & 4 are shifted upwards in the presence of the line probe in windows 1 & 2 respectively. In other words, interpolation settings are most influenced by the nearest visible line segment. Although the task in Part II is the same as that in Part I, i.e. interpolate between the inducers, it is as if the observers instead interpolated between the intervening segment and the inducer local to the location of interpolation on these trials. This result corroborates earlier work using dot-sampled contours (Hon, Maloney, & Landy, 1997; Feldman, 1997; Warren, Maloney, & Landy, 2004), which provides evidence that interpolation mechanisms are strongly influenced by local sources of information, with influence dropping off with increasing distance from the point of interest (see also Singh & Fulvio (2005; 2007) for a similar effect in the context of extrapolation of smooth contours). Of particular note in studies involving dot-sampled contours was the finding of a preference for interpolations that minimized the variance of the angles between two line segments defined by the sampled points (Feldman, 1997; Pizlo, Salach-Golyska, & Rosenfeld, 1997; Warren et al., 2004). Since the large-offset inducers in the current study require an inflecting-contour to globally connect them, an unfavorable condition for visual interpolation (e.g. Kellman & Shipley, 1991; Takeichi et al., 1995; Fulvio et al., 2008), interpolating between the intervening segment and one of the inducers effectively simplifies the observers' task on each trial by allowing for shorter, non-inflecting interpolations. The trade-off, however, results in an inconsistent interpolation (which may nevertheless be the observers' best strategy in carrying out this particular task since they are not penalized for inconsistent performance).

Given that the representation of the interpolating contour is different in the presence of the line probe in Part II, we do not expect an improvement in the precision of the settings with this information since it is apparently inconsistent with the original interpolation. The F-tests, however, demonstrated that although the mean settings in Part II change with the added line probe, there is nevertheless an improvement in the precision of the settings, as 24/32 tests were significant for increases in precision for position settings and 15/32 tests were significant for increases in setting precision for orientation settings (with only 2/32 tests being significant for decreases in setting precision)—a pattern nearly identical to the results for the symmetric inducers.

Small-offset inducers.

The small-offset inducers comprise an intermediate condition between the symmetric and large-offset inducers—they are offset enough to be clearly asymmetric, but are nevertheless relatable and do not require an inflecting-contour for their interpolation. The means and standard deviations of the settings for the small-offset inducers are shown in Figure 7: As in Figures 4 & 6, Part I settings are shown in red and Part II settings are shown in blue.

Figure 7 about here

Across all observers, there is a visible deviation between the Part I and Part II settings, although it is not as large as in the large-offset inducer data. The results of the *t*-tests on the differences between mean settings from Part I to Part II nevertheless yielded 23/32 tests for changes in position significant at the $\alpha = 0.05$ level and 10/32 tests for changes in orientation were significant. As with the large-offset inducer data, there is a definite, consistent pattern in the deviations observed in Part II for the small-offset inducers. The settings in windows 1 & 2 are shifted downwards in the presence of the line probe in windows 3 & 4, and the settings in windows 3 & 4 are shifted upwards

in the presence of the line probe in windows 1 & 2. Both of these shifts are again towards the nearer inducer.

It should be noted that, although, according to the binary formulation of the relatability criteria the small offset inducers are relatable, in graded terms they are clearly less relatable than the symmetric inducers (the point of intersection is much closer to one point of occlusion than the other²). Similarly, from the point of view of the co-circularity model, the small-offset inducers are clearly not cocircular (whereas the symmetric inducers are). This gradedness is clearly reflected in the consistency results: the degree of consistency for the small-offset inducers is clearly lower than for the symmetric inducers, but clearly higher than that of the large offset inducers (non-relatable, and even greater deviation from cocircularity).

Since the results for the small-offset inducers are similar to those of the large-offset inducers, we expect that the results of the precision analysis will also be similar such that the precision of the settings in Part II increases despite the change in the mean settings. Instead, the F-tests do not conform to this expectation as only 5/32 tests were significant for increase in setting precision for position settings (with 1/32 tests being significant for decrease in setting precision) and 4/32 tests were significant for increase in setting precision for orientation settings (with 3/32 tests being significant for decrease in setting precision)—a pattern unlike that of the symmetric and large-offset inducers (See Table 1). Thus, there is also an inconsistency in the interpolation settings between small-offset inducers but unlike the symmetric and large-offset inducers, the settings between small-offset inducers do not become more precise with additional information.

² As one keeps on increasing the offset, eventually the linear extension of one inducer will intersect the other inducer itself, which would then violate the binary formulation of this relatability criterion (as is indeed the case for the large-offset inducer condition).

We considered the possibility that our results depend on our use of the mean (rather than some other measure of central tendency such as the median) in our analyses³. This could occur if, for example, the distributions of first and second settings are very different. We repeated the analyses just described with medians in place of means with no change in conclusions⁴. In any case, setting variability is small compared to the changes in mean (or median) setting that are our measure of inconsistency. Furthermore, we plotted histograms of the z-scores of all of the observers' data, which reveal no evidence of distributional change (see Figure 8). We emphasize that tests of consistency should typically include assessment of possible multi-modality in distributions of settings.

Figure 8 about here

Discussion.

The results of Experiment 1 indicate that symmetric, co-circular inducers evoke single, stable interpolation whereas large-offset inducers do not. Nevertheless, in both cases, adding one of the observer's own settings to the display improved localization of the contour in both conditions. In the third, intermediate, small-offset condition, the degree of internal consistency is relatively low, which points to the importance of co-circularity in contour interpolation. Also, in this condition, the added setting does not lead to an increase in precision.

Recall that Fulvio et al (2008) considered two candidate measures based on interpolation performance. The first was a measure of consistency similar in spirit to the consistency measure considered here and the second was based on the standard

³ We would like to thank one of the reviewers for pointing this out.

⁴ The significance tests revealed 28/32 consistent (i.e. failed to reject consistency) for the no-offset condition, 1/32 consistent for the large-offset condition, and 6/32 consistent for the small-offset condition as compared to 30/32, 4/32, & 9/32 consistent for the respective conditions using the mean.

deviation of interpolation settings. Fulvio et al refer to the reciprocal of setting standard deviation of settings as precision and we use the same terminology here. The motivation for proposing a link between precision and contour strength is evident. If a contour is very “strong” we expect that interpolation settings will vary little (high precision) and if it is “weak” we expect that interpolation settings will vary considerably (low precision).

In the current experiment, all subjects’ position settings exhibited a decline in precision from the symmetric to large-offset inducer conditions for the Part I settings, consistent with the Fulvio et al findings and the expectations of contour strength as determined by degree of cocircularity or relatability. The precision of the Part I orientation settings as a function of inducer offset was highly variable across subjects. Taken together, these results suggest that inducer geometry indeed influences the strength of the interpolating contour by affecting localization at any particular position along the contour.

One puzzling result, however, was the non-systematic benefit in setting precision in the Part II settings with the added line probe. Specifically, the settings for the symmetric and large-offset conditions were more precise with the added probe, while the settings for the small-offset condition were not. We offer the following conjecture as a possible explanation.

In the symmetric inducer condition, the contour is “strong” to begin with (evidenced by the high precision and consistency in subjects’ performance) and the setting in Part II with the additional piece of information should be less variable: there is more consistent information available to estimate the location of the occluded contour. The visual system adopts a strategy that simply combines the available visual information.

In the large-offset inducer case, the contour is very “weak” to begin with, evidenced by the low precision and inconsistency in subjects’ performance. The stimulus

in Part II of the experiment, therefore consists of three segments (the two inducers and the added segment corresponding to the setting from Part I) that are classified as inconsistent with one another by the visual system. The visual system, confronted with inconsistent evidence, may simply discard one of the inducers and interpolate the added segment and the remaining inducer. There is no reason why precision should not be as high as in the reliable case: the added segment is typically reliable to the remaining inducer and the gap over which the visual system interpolates is smaller.

But in the small-offset inducer case the visual system is in a quandary: it is unclear which strategy to follow. If the visual system selects different strategies on each trial, and if the interpolation settings resulting from the two strategies are different, mixing strategies will lead to a decrease or at least no increase in setting precision. Indeed, we find no benefit with the added probe for these stimuli, consistent with the conjecture we propose.

Nevertheless, the results of Experiment 1 provide a validation of the two-point probe test as a test of consistency and a measure of strength of amodal completion: Observers' settings were consistent in the symmetric condition when the inducing contours were reliable and co-circular, but inconsistent when in the large-offset condition when they were non-reliable as expected. The results also suggest that visual interpolation—specifically amodal completion in our stimuli—breaks down in a graded manner with increase in relative offset between the inducers, not in a binary fashion as articulated in the original formulation of the reliability criteria.

EXPERIMENT 2

Experiment 2 applies a version of the two-point probe test developed above to the configuration that typically evokes the Poggendorff illusion. We ask whether the visual system interpolates a stable, self-consistent contour in this configuration. Given the nature of the illusion—a misperception of alignment between the two collinear inducers—it seems unlikely that visual system generates a stable, self-consistent representation of a contour in the occluded portion between the inducers. However, such a hypothesis has been difficult to test previously because of the absence of an experimental criterion for determining whether consistent contour interpolation is taking place. Moreover, contour interpolation has generally not been a component of theories of the Poggendorff illusion. The specific question we address in Experiment 2 is whether observers' interpolation settings of position and orientation in the occluded portions of the Poggendorff configuration are mutually consistent.

Methods.

Observers.

Seven new observers completed two experimental parts at New York University. None were aware of the purpose of the experiment. All had normal or corrected-to-normal vision.

Stimuli.

Part I: oblique line segment extrapolation

Examples of the stimulus displays for Part I are shown in Figure 9A. Each display contained an inducer oriented at 45° and an adjustable line probe. Both were white, 6.05 degrees of visual angle (DVA) long, and 5.8 arcmin thick, with anti-aliasing at the resolution of one-fourth of a pixel.

Figure 9 about here

The inducer pointed upward in half of the experimental trials, and downward in the other half. On any given trial, the inducer could appear either on the left or the right side of an occlusion region defined by a gray rectangle with rounded corners, having width w equal to 6.05 DVA and height equal to 16.94 DVA. The adjustable line probe was always initially presented at a horizontal orientation at the midpoint of the vertical height of the occluder along the edge opposite the inducer. The inducer was given a random height along the edge so that observers could not simply remember their settings across trials. The observers adjusted the position and orientation of the line probe so that it extended (extrapolated) the inducer—i.e. they adjusted the line probe so that it was perceptually collinear with the inducer. In the position-adjustment mode, the probe moved vertically along the edge of the occluder; in the orientation-adjustment mode, it pivoted about the point of occlusion. 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 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.

Part II: line segment interpolation

The Part II trials were much like those in Part I (single window interpolation) of Experiment 1 above (examples of displays are shown in Figure 9A,B). In this part of the experiment, each observer interpolated two types of inducer pairs: (i) two collinear oblique segments (the standard Poggendorff configuration, denoted PC); and (ii) the

configuration of segments that the observer indicated were collinear, based on that observer's mean settings in Part I of the current experiment (the extrapolated collinear configuration, denoted EC). All details concerning the occlusion region, spacing of interpolation windows, and the line probe and its adjustments were identical to those of Part I in Experiment 1 above.

Software and apparatus.

The software and apparatus used in this experiment were the same as those used in Experiment 1.

Procedure.

The task of the observers in both experimental parts was to adjust the position, h , and orientation, ϕ , of the straight-line probe until the combination of settings optimized the percept of collinearity with the single inducer (Part I) or of smooth continuation of the partly-occluded contour defined by the two inducing contours (Part II). The adjustments were made using the same procedure as in Experiment 1.

Design.

In Part I, the inducing line segment could appear on the left or right edge of the occluder and it could point upwards or downwards. Thus, the design was 2 x 2 or 4 trial types, times 10 repetitions for a total of 40 trials in Part I. Each trial obtained paired adjustments of position and orientation of the adjustable line probe to be collinear with the inducer.

The design for Part II contained two inducer pair types (PC and EC; see above) and four window locations. Just as in Experiment 1, one of the two inducers was designated the reference inducer. On any given trial, the reference inducer could appear either on the left or on the right side of the occluder. The slope of the inducers was also randomly designated positive or negative, with counterbalancing occurring over all trials.

Each session thus contained $2 \times 4 \times 2 \times 2$ or 32 trial types, times three repetitions of each combination for a total of 96 trials (48 trials for each inducer pair type—12 per window). Here, each trial obtained paired adjustments of position and orientation of the adjustable line probe to optimize the percept of a smoothly continuing contour between the inducers. Each of the two experimental parts was preceded by a practice session to ensure the observer's understanding of the task and comfort with the controls.

Results and analysis.

In Part I, observers extrapolated a straight-line inducer and set an adjustable line probe to be collinear with the inducer at the opposite edge of a rectangular occluder. In Part II, observers interpolated the line segments in the standard Poggendorff configuration (PC condition), and the line segments consistent with those they indicated as collinear in their Part I settings (EC condition). As in Experiment 1, based on preliminary analysis we were able to combine conditions that differed only in left-right reflection and upwards-downwards inducer presentation. Table 2 contains the results of the statistical tests described below.

Extrapolation of an oblique line segment

In Part I, we recorded observers' settings relative to what would be the actual linear extrapolation of the inducer. All 7 observers set the probe to be vertically-offset below the true extrapolation (see Figure 10). On average, this offset was 1.61 DVA (depicted by the dashed black line in Figure 10), which is highly significant ($t(6) = 16.2$, $p < 0.01$). This finding is consistent with the Poggendorff illusion: the illusion is characterized by the percept of the inducers being vertically offset across the gap and hence non-collinear, so the undershooting appears to counteract the perceived vertical offset, allowing them to appear collinear. In their orientation settings, observers exhibited on average a small but significant deviation of 2.4 degrees less than the linear extrapolation orientation of 45 degrees ($t(6) = -3.09$, $p = 0.02$).

Figure 10 about here

Interpolation of the Extrapolated Collinear configuration (EC)

In Part II, the observers interpolated two inducer pair types, the first of which was the configuration that they indicated is collinear in Part I (EC). For this condition, we presented the original reference inducer paired with an inducer determined by the observer's mean position and orientation settings from Part I. For all seven observers, this particular configuration was not actually collinear even though they extrapolated it as such in Part I. These inducer pairs thus require an inflecting contour to smoothly interpolate them—a property which has been shown to be detrimental to interpolation (e.g. Kellman & Shipley, 1991; Takeichi et al., 1995; Fulvio, et al., 2008). Thus, for this condition, we are interested in how the discrepancy between the subjective impression of collinearity and the physical non-collinearity impacts the consistency of the visually interpolated contour between these inducer pairs.

Figure 11 about here

The means and standard deviations of the position and orientation settings are plotted for each of the seven observers in Figure 11. Observers make position settings that are consistent with the straight line that connects the end points of the inducers. Therefore, even though observers set these inducers as collinear in Part I, all agree that they are not collinear in Part II. Visual inspection of the data suggests that the orientation settings conform to a different percept than the position settings as they are consistently steeper than the straight-line join of the two inducers' endpoints. To verify this dissociation quantitatively, we performed t-tests (7 observers x 4 interpolation windows =

28 tests) on the difference between each observer's mean orientation settings in the EC condition and the orientation of the straight-line join. The results of these tests revealed that 28 out of 28 were significantly different, thereby quantitatively confirming the qualitative trend. In a second set of t-tests, we examined the difference between each observer's mean orientation settings in each of the four windows in the EC condition and 45 degrees. The 45 degree orientation corresponds to that of the straight-line join in the PC condition, which we note is larger than the straight-line join in the EC condition. Thus, this test was motivated by the possibility that observers' orientation settings were subject to a Poggendorff illusion-like effect even for the EC inducers, which in effect causes the inducers to appear more offset than they are, and hence have a steeper intervening orientation than they do. The results of these tests revealed that 16 out of 28 were not significantly different than 45 degrees, with the remaining 12 cases that were significantly different having orientations exceeding even the 45 degrees.

We conclude then, that the position and orientation settings correspond to different perceptual interpretations of the relationship between the line segments in the EC condition. Although not the primary purpose of this study, this result is nevertheless interesting in that it hints that the subjective impression of collinearity indicated by the Part I extrapolation settings may have some bearing on the orientation settings, but not the position settings, when this configuration is interpolated. Future work devoted to investigating interpolation between linear inducers may provide further insight.

Interpolation of the standard Poggendorff configuration (PC)

The second inducer pair type observers interpolated in Part II was the standard Poggendorff configuration (PC). Physically, the line segments in this configuration can be interpolated linearly using a straight line with a 45° orientation. However, the

subjective impression typically evoked by this stimulus is that these line segments are not collinear, and instead are offset (and thus would require an inflecting contour to be interpolated smoothly).

Figure 12 about here

The means and standard deviations of the position and orientation settings for the PC condition are plotted in Figure 12 for all seven observers. Just as in the EC condition, observers made their position settings consistent with the straight line join that connects the endpoints of the inducers. Recall that in the EC condition, we could not find evidence for a stable representation because the position and orientation settings did not conform to the same interpolating-contour. Specifically, the position settings agreed with the straight-line connection between the inducers' endpoints, whereas the orientation settings were much steeper, tending to agree with the orientation of the straight-line connection between the inducers' endpoints in the PC condition. Focusing our attention now on the orientation settings in the PC condition, we will look for similar evidence that would suggest dissociation between the two setting types in this configuration.

Inspection of the orientation data for the PC condition reveals a trend like that in the EC condition, as the orientation settings appear to be considerably steeper than that of the straight line defined by the position settings. We performed t-tests on the mean difference between the orientation settings in each window and 45 degrees, which is consistent with the orientation of the line defined by the position settings. These tests resulted in a rejection of 26 out of 28 for 45 degrees—evidence for a dissociation between the percepts defined by the position and orientation settings. The inducers in the PC configuration therefore also do not evoke a single, stable representation as predicted.

Figure 13 about here

It is interesting that observers' position settings do in fact conform to the straight line that actually interpolates the inducers in the PC condition—which would not be predicted by previous theories of the Poggendorff. Although we cannot explain the Poggendorff illusion with these data, we can rule out one processing strategy that may have led to the above results: that observers are simply averaging the influence of two extrapolated Poggendorff effects⁵. To do this, we quantified each observer's "Poggendorff effect" by the average deviation in position and orientation settings from the correct extrapolation of the oblique segment in Part I. We then computed the two individual inducer extrapolations the observers should have made in the interpolation condition of Part II given their respective Poggendorff effects (see Figure 13 for an example). At each window location, we averaged the positions of the two extrapolated Poggendorff effects to yield the settings that observers should have made had they adopted this strategy during the interpolation task. We then performed t-tests (7 observers x 4 windows = 28 tests), first on the difference between each observer's Part II mean position settings and the position of the actual straight-line connection between the inducers, and second on the difference between each observer's Part II mean position settings and the position of the line resulting from the average of two Poggendorff effects. We could not reject any of the 28 tests for the straight-line join at the $\alpha = 0.05$ level, but we rejected 16 of 28 tests for the averaged Poggendorff effects. (Note that the 12 cases that we could not reject were in the central window locations where the predictions of the average Poggendorff effects and the true straight line are

⁵ We would like to thank Wilson S. Geisler for recommending this analysis in a personal communication.

close). Thus, as with the EC condition, the subjective impression of non-collinearity between PC inducers indicated by the Part I extrapolation settings may have some bearing on the orientation settings, but not the position settings, when this configuration is interpolated.

Discussion.

Our investigation of contour completion in the standard Poggendorff configuration shows that Poggendorff illusion-like effects appear in the extrapolation of an oblique line segment across a rectangular gap and in the orientation judgments of the interpolation between the collinear line segments in the standard Poggendorff configuration whereas position settings conform to the straight-line join between the inducers' endpoints. Thus, dissociations between position and orientation processing within the gap of the Poggendorff configuration exist and as predicted, lead to inconsistencies in observer judgments.

GENERAL DISCUSSION

When making sense of scenes in our environment, whether to act on them or simply to passively view them, we are almost always faced with the task of visually-completing portions of objects that are partly hidden. A large body of research has provided important insights into this process and how it affects performance on a wide range of visual tasks. With recent exceptions including Fulvio et al (2008), however, the question as to whether “completed” representations are actually complete and stable has remained unaddressed.

In the current study, we put forth a new experimental criterion for performance reflecting complete and stable contour interpolation: self-consistency. In order to measure performance according to this criterion, we proposed the two-probe interpolation method. In Experiment 1, which was designed to test the validity of the method and criterion, we provided evidence that visual interpolation judgments for a given pair of inducers are not always consistent with a single, stable, smooth contour. In particular, the inconsistencies we found occurred for those conditions that would not be expected to produce “strong” representations according to past findings in the literature. This offers support for the method and criterion as effective tools in probing the representations derived by the visual system.

In Experiment 2, we tested for self-consistent amodal completion in a context where the inducing contours are perceived as misaligned despite being collinear—namely, the Poggendorff configuration. We found that observers’ judgments of position and orientation behind the rectangular occluder are mutually inconsistent. Again, we find inconsistency in a condition that is not expected to produce a complete, stable representation. The steep orientation settings made between the linear inducers in the Poggendorff configuration are consistent with past descriptions of the Poggendorff illusion as arising from the misestimation of the orientation of the virtual line connecting

the endpoints of the linear inducers as being too steep (e.g. Weintraub & Tong, 1974; Day & Dickinson, 1976; Morgan, 1999). On the other hand, the position settings are consistent with the straight-line connection between the two inducers' endpoints. Thus, in both experiments, the two-probe method and self-consistency criterion yielded performance and characterizations consistent with past findings in the literature, thereby providing support for their validity.

In Experiment 1 we varied the relatability of linear inducers separated by an occluding region and found close agreement between relatability (interpreted in graded terms) and consistency as measured by the task we developed. We emphasize that we do not claim that this relation would hold only for partially-occluded contours. Our test of consistency could readily be adapted to other stimulus conditions requiring visual interpolation, including the sampled contours of Field, Hayes, and Hess (1993) and of Warren, Maloney & Landy (2002) where there is no occluding region.

Revisiting the isomorphism hypothesis, it is easy to believe that the visual representation of the world around us is a mirror of the environment outside us even if, on occasion, we might acknowledge that the mirror is distorted. The contribution of Shepard in formulating the isomorphism hypothesis was to make precise what it means for a representation to “mirror” the world. In the world every object has its fixed place and the locations and orientations that are the focus of corresponding psychophysical judgments are consistent. The isomorphism hypothesis entails that the visual representation has the same consistency.

Like several previous studies, our study calls this naïve belief in the “isomorphism hypothesis” into question (e.g., Wenderoth, 1983; Gillam & Chambers, 1985; Smeets et al, 2002). The many judgments we may make about a single occluded contour are not always consistent, and, for the Poggendorff configuration, these inconsistencies are readily interpretable. Indeed, there are alternative ways to

characterize internal representations and the role they play in perception and action. Maloney (2002), for example, argued that any coherent interpretation of visual perception as Bayesian inference required that sensory information encoded as probability distributions be combined with arbitrary loss functions without an intermediate “pictorial” representation of scenes. Recent work in visual cue combination (Landy et al, 1995) implies that the visual information available to influence behavior does include estimate of visual uncertainty as well as estimates of position, etc. Graf, Warren, & Maloney (2005) found that observers had access not simply to estimates of the extrapolated position of a moving object that passed beneath an occluding contour but also to estimates of the uncertainty of the position estimates. Shah & Singh (2007) found that observers could combine their implicit knowledge of their variance in extrapolating curved motion paths with an externally specified loss function in a near-optimal fashion. Estimating position and orientation are two distinct tasks, and could plausibly involve different loss functions; so there is no reason to suppose that the two sets of judgments should be mutually consistent unless visual processing constrains them to be so (Fulvio et al., 2008).

The outcomes of the experiments reported here are not consistent with the claim that the visual representation is simply a “picture”. Previous work in cue combination provides strong evidence that initial visual information (“cues”) about properties of the scene such as depth is paired with measures of the variance of each cue (Landy, Maloney, Johnston & Young, 1995). Bayesian approaches likewise assume that visual processing has access to estimates of the distribution of visual properties (Knill & Richards, 1996).

Following Graf et al (2005), we conjecture that, associated with each estimate of a property of the scene, there are additional estimates of parameters conveying distributional information about the original estimate. These could include an estimate of

the variance or standard deviation of the original estimate and possibly more. The visual representation then is effectively a collection of information about the scene, but not a reproduction of the scene. The added distributional information tells us, for example, not just what is in the scene but also how reliable our estimates of scene properties are.

Given a particular visual task, we combine relevant information and make a judgment. There is no guarantee that the outcome of these judgments add up to a single, stable world.

CONCLUSION

The primary purpose of this paper was to define and test an experimental criterion, namely, consistency, as a measure of the stability of visual interpolation in amodal completion. In Experiment 1, we found that observers are consistent by this criterion for some pairs of inducing contours and inconsistent for others, and that consistency is in agreement with cocircularity and relatability — as long as these are interpreted in graded terms (see also Fulvio, et al., 2008). In Experiment 2, we found, as expected, an inconsistency in amodal completion of the Poggendorff configuration—specifically a dissociation between observers' settings of interpolation position and orientation. Both sets of results validate the two-probe method as an experimental test of whether stable, self-consistent, contour interpolation is taking place in any given configuration.

Researchers sometimes describe the “strength”, “goodness”, or “salience” of partly-occluded contours. We cannot, of course, prove that our measure of consistency coincides with “strength”, “goodness” or “salience” simply because there are no agreed-upon experimental procedures to measure “strength”, etc. We could, of course, ask observers to rate the “strength”, etc of contours as has already been done by other researchers (e.g. Kellman & Shipley, 1991; Takeichi et al., 1995). However, we

emphasize that our measure of consistency is based on the observer's interpolation performance while such rating measures require that the observer understand what is meant by "strength", etc. Moreover, even if we know that an observer has rated a contour as "strong" or "weak", we do not know what implications this rating has for observers' interpolation performance. Consequently, we regard the proposed consistency measure as an operationalization (Bridgman, 1927) of "strength" of partly-occluded contours (and of a contour's representation in general). We are fully prepared for the possibility that future research will demonstrate that our measure must be supplemented by other measures that capture other aspects of the intuition of "strength". We ask only that these measures be based on what observers do and can do in localizing the occluded portions of partly-occluded contours.

ACKNOWLEDGMENTS

This research was funded in part by a Graduate Research Fellowship from the National Science Foundation (JMF), Grant CCF-0541185 from the National Science Foundation (MS), and Grant EY08266 from the National Institute of Health (LTM). We would like to thank two anonymous reviewers for their helpful comments and suggestions on a previous version.

REFERENCES

- Anderson, B.L., Singh, M., & Fleming, R. (2002). The interpolation of object and surface structure. *Cognitive Psychology*, *44*, 148-190.
- Boyaci, H., Doerschner, K., & Maloney, L.T. (2006). Cues to an equivalent lighting model. *Journal of Vision*, *6*(2):2, 106-118, <http://journalofvision.org/6/2/2/>, doi:10.1167/6.2.2.
- Brainard, D.H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433-436.
- Bridgman, P. (1927). *The Logic of Modern Physics*. New York: MacMillan.
- 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.
- Day, R.H. & Dickinson, R.G. (1976). The components of the Poggendorff illusion. *British Journal of Psychology*, *67*(4), 537-552.
- Ernst, M.O., & Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*, 429-433.
- Fantoni, C. & Gerbino, W. (2003). Contour interpolation by vector-field combination. *Journal of Vision*, *3*(4), 281-303, <http://www.journalofvision.org/5/8/341/>, doi:10.1167/5.8.341.
- Feldman, J. (1997). Curvilinearity, covariance, and regularity in perceptual groups. *Vision Research*, *37*(20), 2835-2848.
- Field, D.J., Hayes, A., & Hess, R.F. (1993). Contour Integration by the Human Visual System: Evidence for a local "association field". *Vision Research*, *33*(2), 173-193.
- Fulvio, J.M. & Singh, M. (2006). Surface geometry influences the shape of illusory contours. *Acta Psychologica*, *123*, 20-40.

- Fulvio, J.M., Singh, M., & Maloney, L.T. (2006). Combining achromatic and chromatic cues to transparency. *Journal of Vision*, 6(8):1, 760-776, <http://journalofvision.org/6/8/1/>, doi:10.1167/6.8.1.
- Fulvio, J.M., Singh, M., & Maloney, L.T. The precision and consistency of visual interpolation. *Vision Research*, 48(6), 831-849..
- Gillam, B. & Chambers, D. (1985). Size and position are incongruous: measurements on the Müller-Lyer figure. *Perception & Psychophysics*, 37(6), 549-556.
- Graf, E.W., Warren, P.A., & Maloney, L.T. (2005). Explicit estimation of visual uncertainty in human motion processing. *Vision Research*, 45, 3050-3059.
- 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., and 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.
- Kanizsa, G. (1979) *Organization in Vision*. Praeger, New York.
- Kellman, P.J. & Shipley, T.F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology*, 23, 141–221.
- Koenderink J.J., van Doorn A.J., Kappers A.M. (1996). Pictorial surface attitude and local depth comparisons. *Perception & Psychophysics*, 58(2):163-73.
- Landy, M.S., Maloney, L.T., Johnston, E.B., & Young, M. (1995). Measurement and modeling of depth cue combination: In defense of weak fusion. *Vision Research*, 35, 389-412.
- 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 Heyer, D. & Mausfeld, R. [Eds], *Perception and the Physical World: Psychological and Philosophical Issues in Perception*. New York: Wiley, pp. 145-189.
- Michotte, A., Thinès, G., & Crabbé, G. (1964). Les compléments amodaux des structures perceptives. Louvain, Belgique: Publications Universitaires de Louvain.
- Morgan, M.J. (1999). The Poggendorff illusion: a bias in the estimation of the orientation of virtual lines by second-stage filters. *Vision Research*, 39, 2361-2380.
- Oruç, I, Maloney, L.T., & Landy, M.S. (2003). Weighted linear cue combination with possibly correlated error, *Vision Research*, 43, 2451-2468.
- Parent, P. & Zucker, S.W. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Transaction 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-

- Shah, S., & Singh, M. (2007). Perceptual decisions under risk in a motion-extrapolation paradigm [Abstract]. *Journal of Vision*, 7(9):614, 614a, <http://journalofvision.org/7/9/614/>, doi:10.1167/7.9.614.
- Shepard, R. (1981). Psychophysical complementarity. In Kubovy, M. & Pomerantz, J.R. [Eds], *Perceptual Organization*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 279-341.
- Shepard, R.N. & Chipman, S. (1970). Second-order isomorphism of internal representations: Shapes of states. *Cognitive Psychology*, 1, 1-17.
- 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(6), 783-798.
- Singh, M. & Hoffman, D.D. (1999). Completing visual contours: The relationship between relatability and minimizing inflections. *Perception & Psychophysics*, 61, 943-951.
- Smeets, J.B., Brenner, E., de Grave, D.D., & Cuijpers, R.H. (2002). Illusions in action : consequences of inconsistent processing of spatial attributes. *Experimental Brain Research*, 147(2), 135-144.
- Stanley, D.A. & Rubin, N. (2003). fMRI activation in response to illusory contours and salient regions in the human lateral occipital complex. *Neuron*, 37, 321-323.
- 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.

Weintraub, D.J., & Tong, L. (1974). Assessing Poggendorff effects via collinearity, perpendicularity, parallelism, and Oppel (distance) experiments. *Perception & Psychophysics*, *16*, 213-221.

Wenderoth, P. (1983). Identical stimuli are judged differently in the orientation and position domains. *Perception & Psychophysics*, *33*(4), 399-402.

TABLES

Table 1		
<i>Summary of Experiment 1 Results.</i>		
Symmetric (non-offset) inducers		
	Consistency	Precision increase
Setting Type		
Position	30/32	22/32
Orientation	31/32	12/32 ¹
Large-offset inducers		
	Consistency	Precision increase
Setting Type		
Position	4/32	24/32
Orientation	24/32	15/32 ²
Small-offset inducers		
	Consistency	Precision increase
Setting Type		
Position	9/32	5/32 ³
Orientation	22/32	4/32 ⁴

¹ 4/32 precision decreases

² 2/32 precision decreases

³ 1/32 precision decreases

⁴ 3/32 precision decreases

Table 1. Summary of Experiment 1 Results.

Table 2		
<i>Summary of Experiment 2 Results.</i>		
Extrapolation		
Deviation from true line		
Setting Type		
Position	-66.62 pixels	
Orientation	-2.4 degrees	
EC inducer interpolation		
	True line orientation	Linear EC orientation
EC orientation settings	16/28	0/28
PC inducer interpolation		
	True line	Average Poggendorff
Setting Type		
Position	28/28	12/28
Orientation	2/28	17/28

Table 2. Summary of Experiment 2 Results.

FIGURE LEGENDS

Figure 1. A case of occlusion. (A) Observers typically interpret the scene as containing a complete dark-colored surface behind a light-colored surface rather than two dark objects abutting a light-colored one. (B) An attempt to picture the observer's visual representation of 1A after completion of the occluded object and its contours.

Figure 2. A linking hypothesis: Measuring the strength of amodal completion. A. Interpolation between two linear inducers. The red solid line marks the locus of first settings where an observer would interpolate the partly-occluded contour defined by the two inducers. B. Interpolation with an added probe. The hypothetical locus of second interpolations (conditional on the added probe) is shown as a dashed blue line. We take the deviations between first and second loci of interpolation ("inconsistency") as a measure of the weakness of the amodal completion of the partly-occluded contour.

Figure 3. Examples of stimulus displays used in Experiment 1. (A) Part I displays: The observer could control the height and orientation of the short linear segment that appears in the vertical slit in the occluded region. (B) Part II displays: same as (A) except a second vertical slit appears with a fixed line segment positioned and oriented as the observer's own mean settings from Part I. See text for details. (C) Three vertical offsets were used in the study: no offset ($\Delta = 0$), small offset ($\Delta = \frac{2}{3} w \tan \theta$), and large offset ($\Delta = \frac{4}{3} w \tan \theta$).

Figure 4. Experiment 1 mean settings for the no offset inducers. 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 settings of Part I

are drawn with a red-blue color scheme, and those of Part II are drawn with a blue-black color scheme.

Figure 5. Schematic depiction of one of the pairs of settings in Figure 4. Left panel:

The observer made the depicted settings in Windows 2 & 4 in Part I. Right panel: When the observer's mean settings for Window 4 were shown (depicted by the red segment), the observer made the setting in Window 2 depicted by the green segment. We see that in comparison to the observer's original settings for Window 2 (depicted by the black segment), the observer is highly-consistent from Part I to Part II.

Figure 6. Experiment 1 mean settings for the large offset inducers. The format is identical to that of Figures 5.

Figure 7. Experiment 1 mean settings for the small offset inducers. The format is identical to that of Figures 5 & 7.

Figure 8. Histograms of the z-scores of all of the observers' data in each of the four windows. Settings for Part I are drawn in red, those for Part II, in blue. The histograms reveal no evidence of bimodality even for the large-offset condition

Figure 9. Examples of stimulus displays used in Experiment 2. (A) Part I displays:

The observer could control the height and orientation of the horizontal segment so that it appears to be collinear with the oblique segment. (B) Part II displays: same as Experiment 1, Part I (recall Figure 3A). Observers interpolated two conditions for Part II: the standard Poggendorff configuration (PC, left) and each observer's own extrapolated collinear configuration from Part I (EC, right for one observer). See text for more details.

Figure 10. Experiment 2 mean settings for extrapolation task. The mean extrapolated position and orientation settings of all eight observers are depicted by the colored segments on the right edge of the rectangle. The average settings

for all observers are depicted by the dashed segment, and the true collinear extrapolation is depicted by the solid segment.

Figure 11. Experiment 2 mean settings for the EC condition of the interpolation task. The means and standard deviations of position settings are shown as points with error bars (± 1 SD). The means of the orientation settings are represented by short line segments whose tangent orientation is the mean setting. The dashed line is the straight-line connection of the two inducer endpoints.

Figure 12. Experiment 2 mean settings for the PC condition of the interpolation task. The format is identical to that of Figure 11.

Figure 13. Depiction of the average Poggendorff effect analysis. The extrapolated Poggendorff effects were first computed (depicted in red) based on each observer's results from Part I. The average at each window location was then taken (depicted in blue) and the observer's mean positions settings were compared to the averages and to the true collinear predictions depicted in black. See text for more details





Figure 1

A



B






Figure 2



A



B

Figure 3AB




Figure 3C

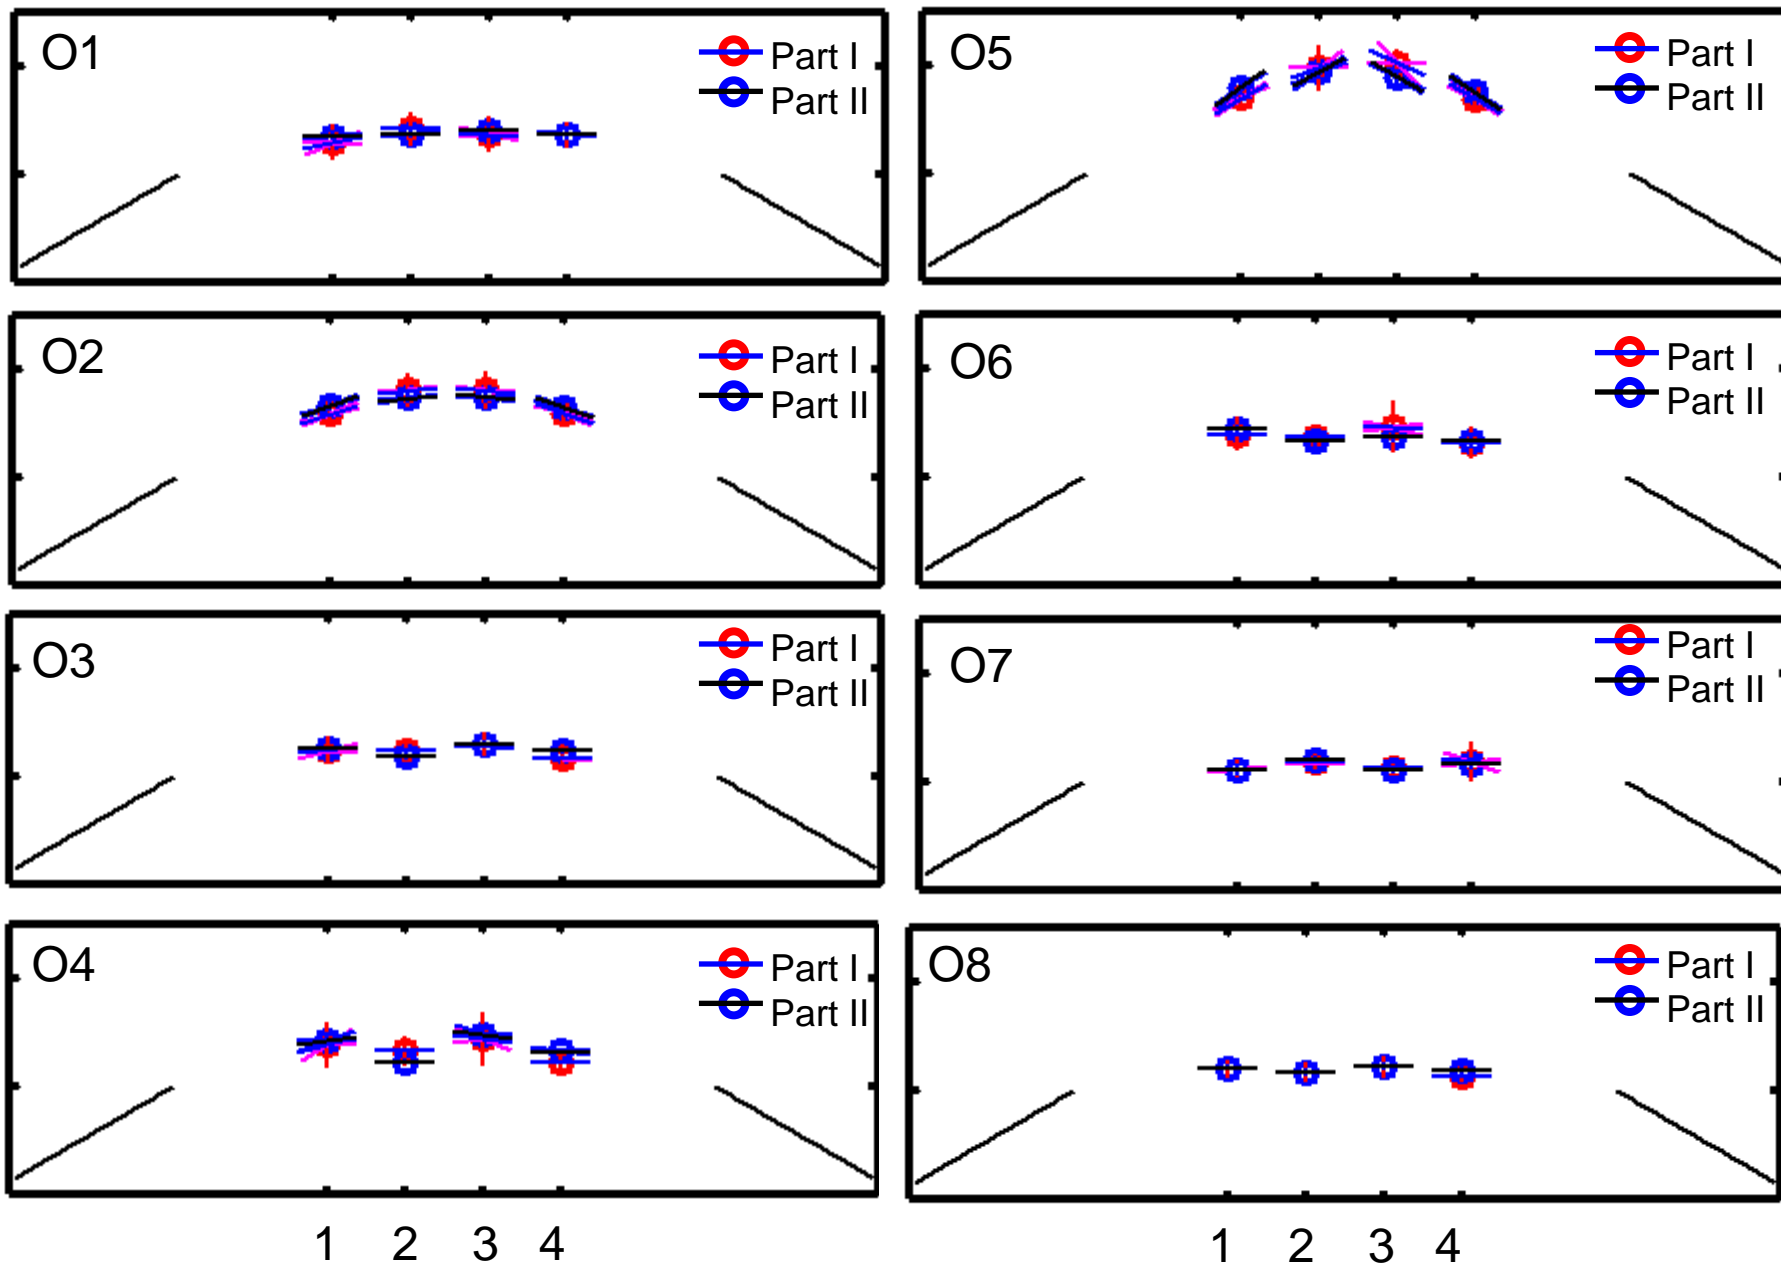
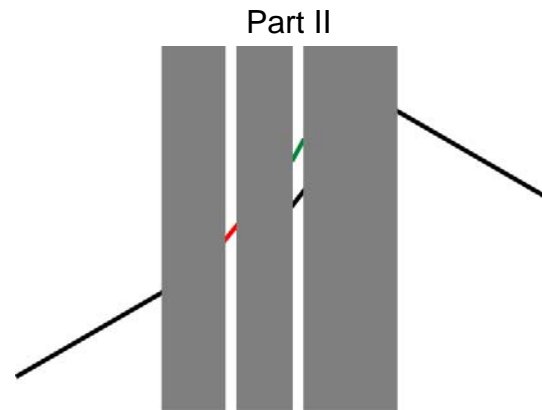





Figure 4

Interpolation Window



A.

B.

Figure 5

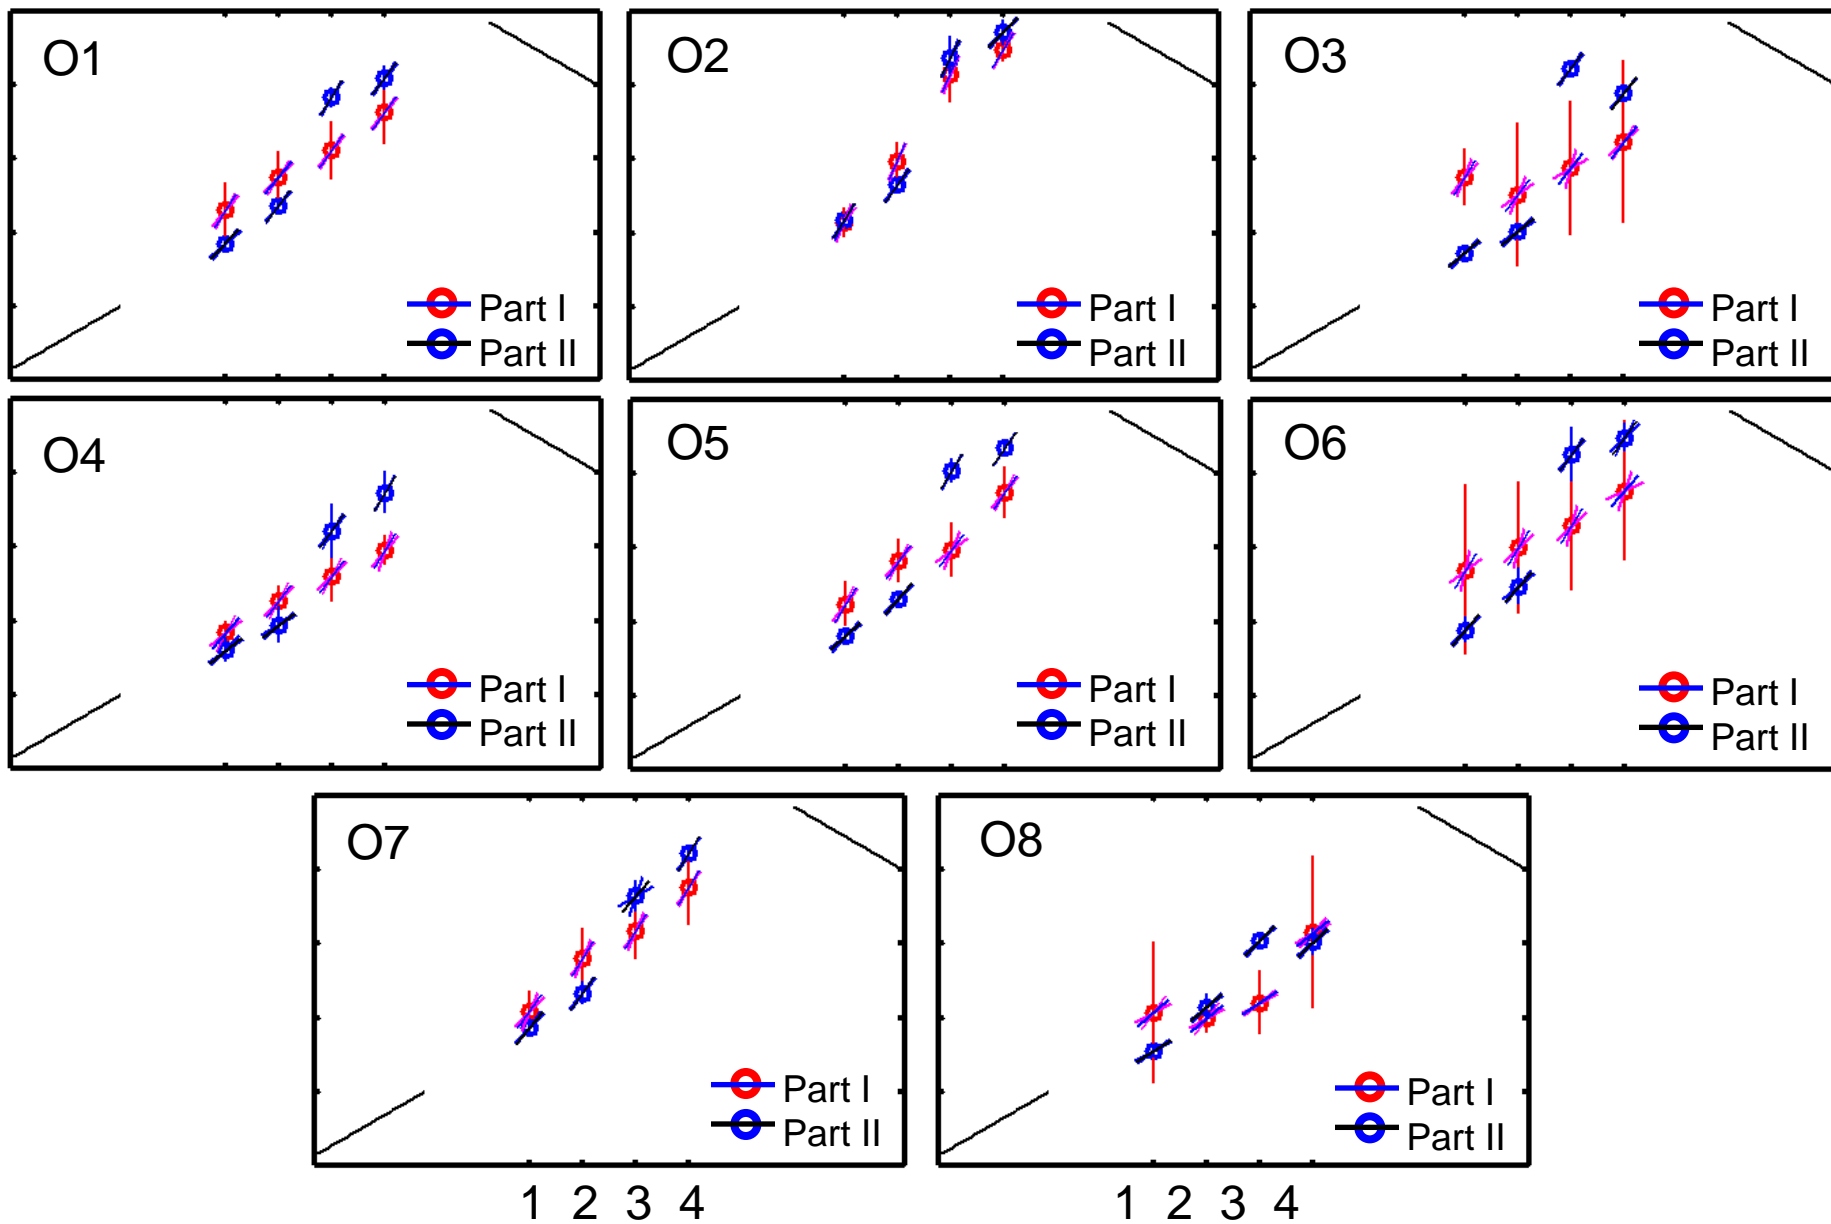


Figure 6

Interpolation Window

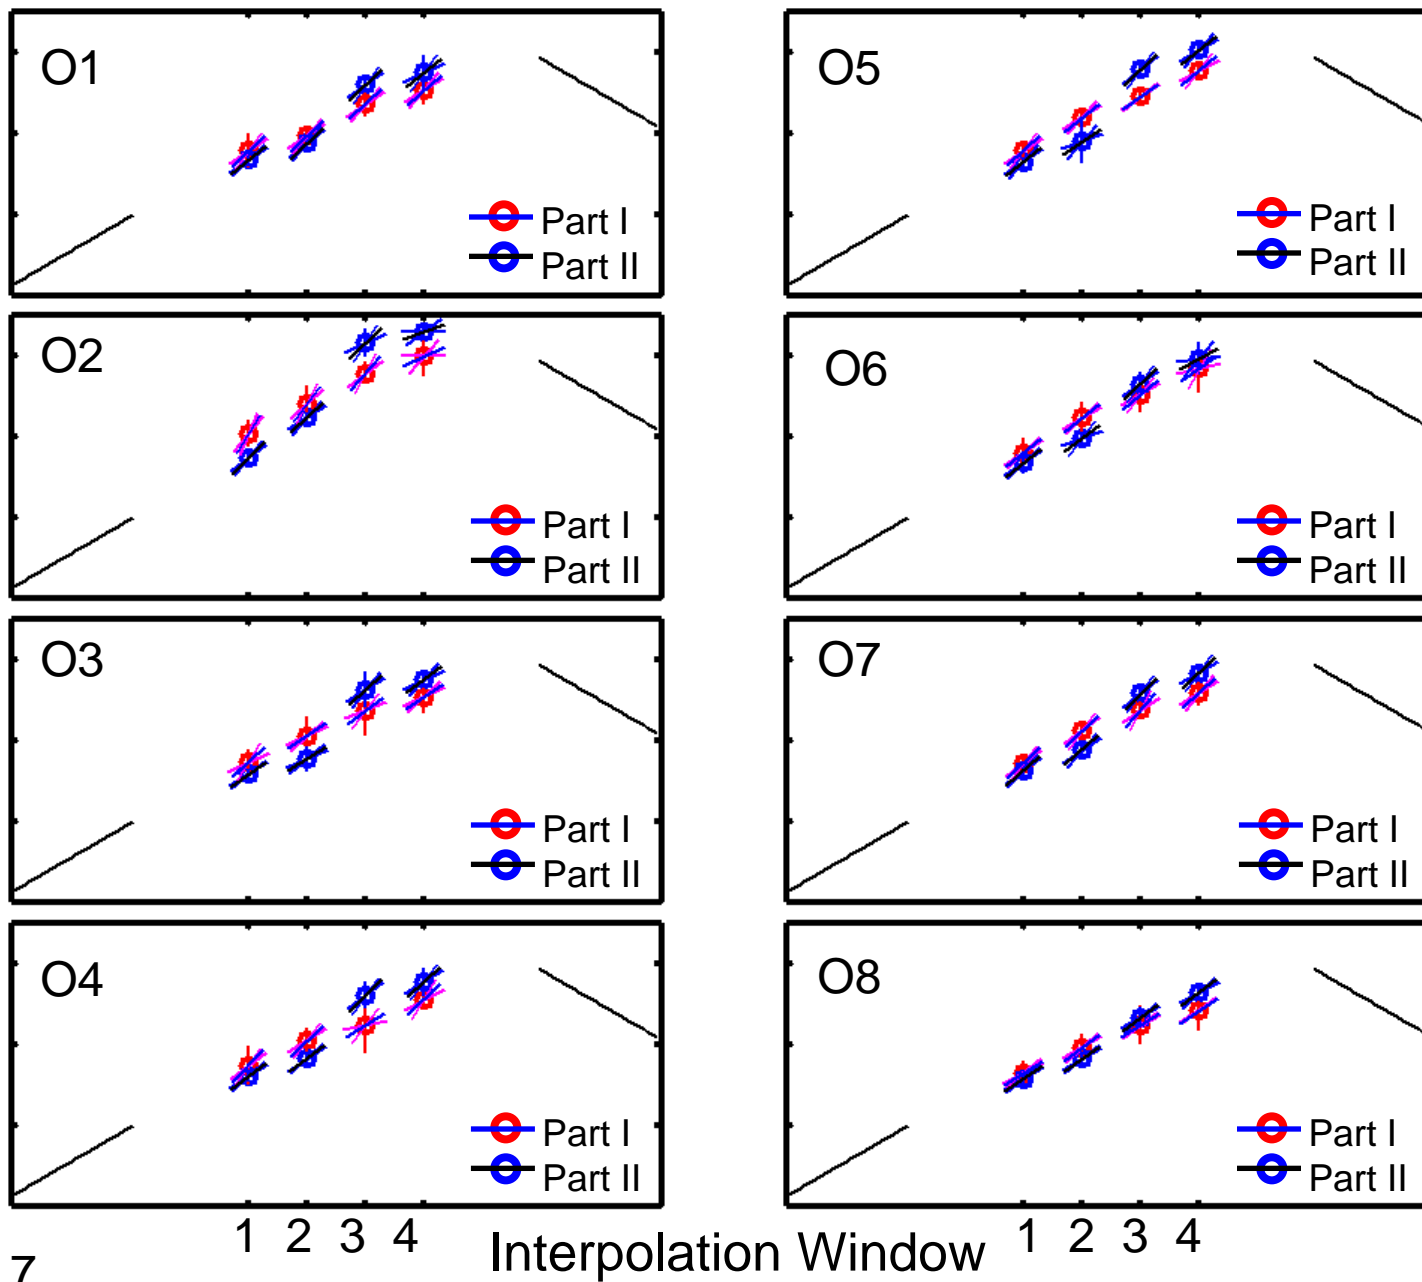


Figure 7

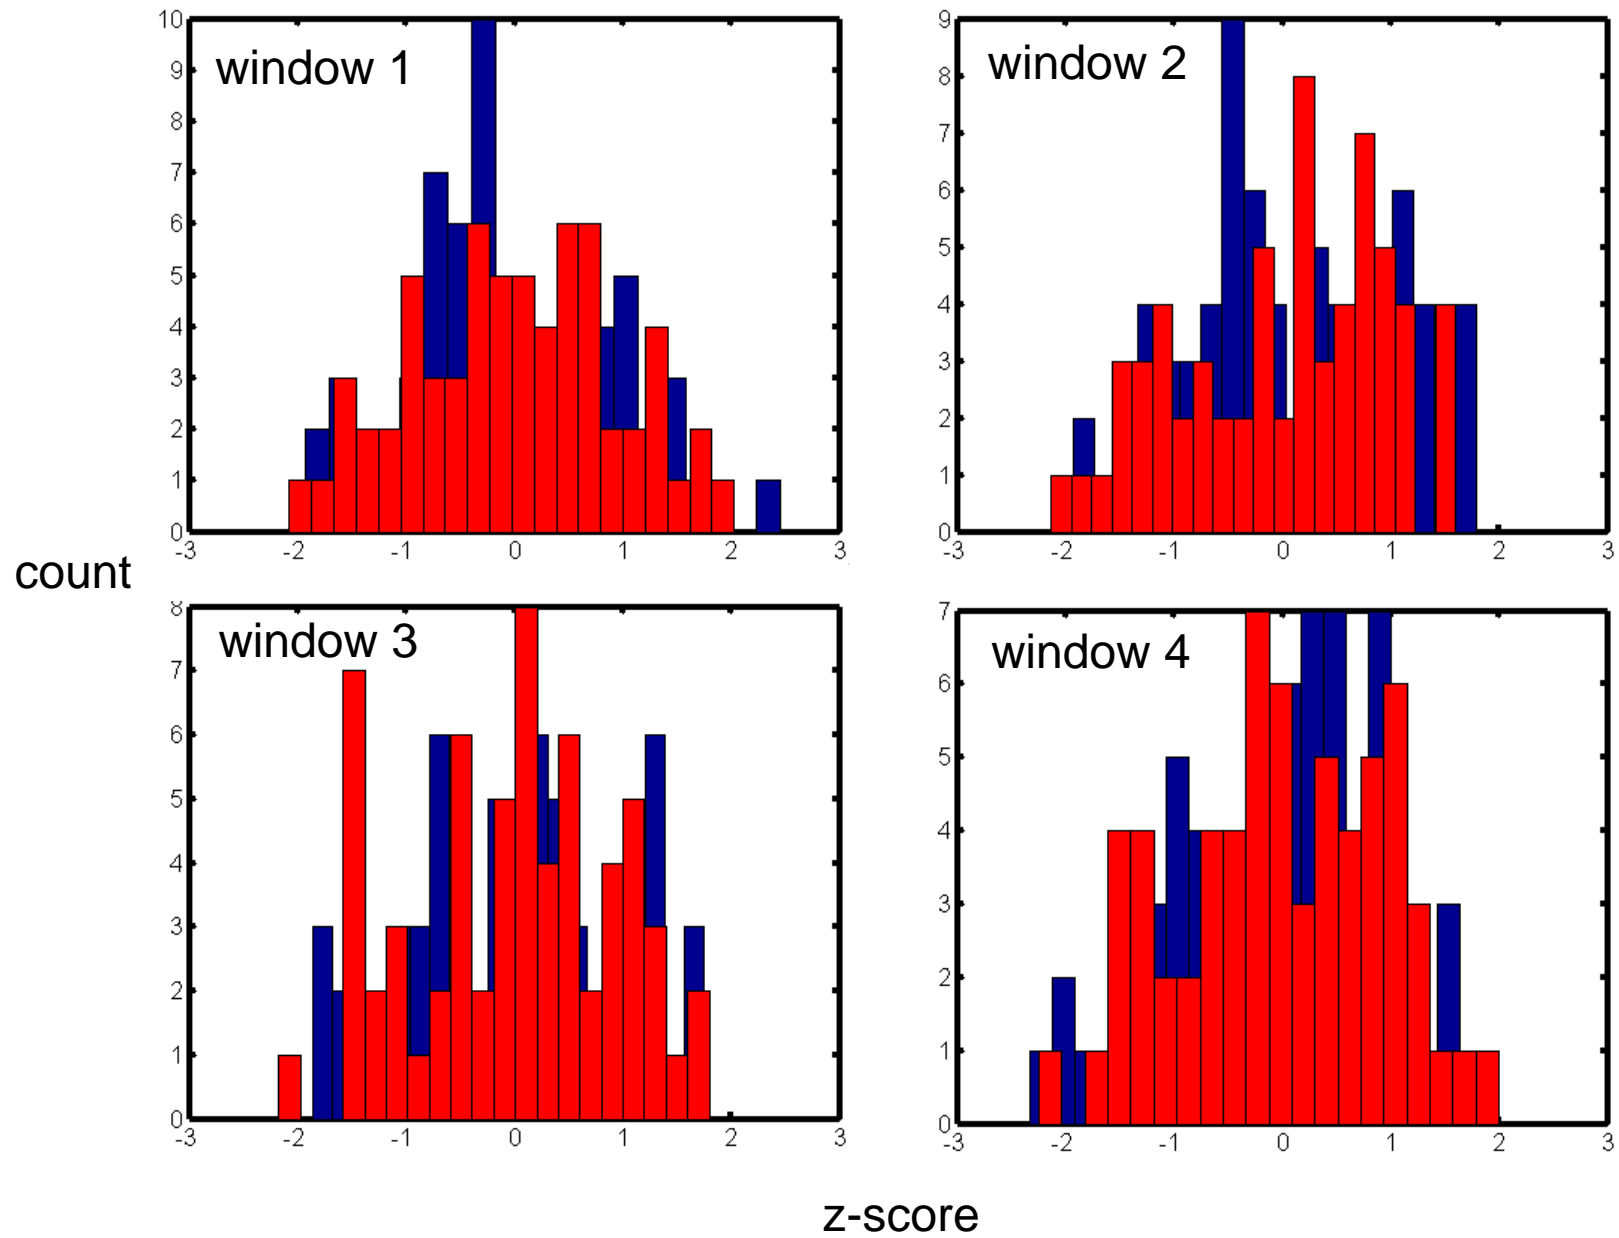




Figure 8



A



B

Figure 9




Figure 10




Figure 11




Figure 12




Figure 13

Table 1		
<i>Summary of Experiment 1 Results.</i>		
Symmetric (non-offset) inducers		
	Consistency	Precision increase
Setting Type		
Position	30/32	22/32
Orientation	31/32	12/32 ¹
Large-offset inducers		
	Consistency	Precision increase
Setting Type		
Position	4/32	24/32
Orientation	24/32	15/32 ²
Small-offset inducers		
	Consistency	Precision increase
Setting Type		
Position	9/32	5/32 ³
Orientation	22/32	4/32 ⁴

¹ 4/32 precision decreases

² 2/32 precision decreases

³ 1/32 precision decreases

⁴ 3/32 precision decreases

Table 1

Table 2		
<i>Summary of Experiment 2 Results.</i>		
	Extrapolation	
	Deviation from true line	
Setting Type		
Position		-66.62 pixels
Orientation		-2.4 degrees
	EC inducer interpolation	
	True line orientation	Linear EC orientation
EC orientation settings	16/28	0/28
	PC inducer interpolation	
	True line	Average Poggendorff
Setting Type		
Position	28/28	12/28
Orientation	2/28	17/28

Table 2