Visual interpolation and extrapolation of contours

Jacqueline M. Fulvio, Manish Singh and Laurence T. Maloney

In Gepshtein, S., Singh, M. & Maloney, L. T. (2015), *Oxford Handbook of Computational Perceptual Organization.* New York: Oxford University Press, *in press.*

1. Introduction

In everyday, cluttered scenes, the visual system faces three problems in identifying contours (Warren, Maloney, & Landy, 2004; Singh & Fulvio, 2007). The first is the *detection problem*: deciding whether one or more contours are present in the scene. The second is the *grouping* or *segmentation problem*: deciding which visual fragment belongs to which contour (Figure 1A). Moreover, many of the contours that we encounter are only partly visible due to occlusion. Even after solving the grouping problem, the visual system still needs to estimate the shape of the contour both where it is visible and where it is not (the *interpolation/extrapolation problem* illustrated in Figure 1b).

There is considerable prior work evaluating human performance in detection and grouping (Grossberg & Mingolla, 1985; Parent & Zucker, 1989; Kellman & Shipley, 1991; Field, Hayes, & Hess, 1993; Elder & Goldberg, 2002; Geisler, Perry, Super, & Gallogly, 2001). In what follows, we focus on human performance in estimating the shape of contours.

Figure 1 about here

The shape problem is difficult in part because it is under-constrained, a form of the problem of induction (Hume, 1748/1910). Any pair of fragments can always be connected – *interpolated* – smoothly by any of a large set of possible segments (Figure 1b). A natural goal of research is therefore to identify the constraints that the visual system imposes in estimating contour shape (e.g., Ullman, 1976; Singh & Fulvio, 2005; 2007). A second goal is to decide whether there are constraints that lead to accurate shape interpolation in the terrestrial environment (e.g., Geisler, Perry, Super & Gallogly, 2001). We would hope that advances toward either goal would give us clues as to how – and why – the visual system estimates contour shapes as it does.

The problem of shape completion is closely related to the often-cited concept of "good continuation." Wertheimer (1923) originally proposed the principle of good continuation as a way of choosing between different possible extensions of a contour segment (see, e.g., his Figures 16-19). The principle suggests that the *ideal* extension of a contour segment is one that continues the generative process that gives rise to it. However a precise mathematical characterization of good continuation by human vision has been elusive. Specific questions include: (i) what geometric properties of the contour does the visual system use in extending its shape—orientation, curvature, rate of change of curvature, etc.? and (ii) how does it combine the contributions of these properties to define the shape of the extended contour?

Three approaches to modeling shape interpolation have been used: (i) to specify

a fixed class of shape primitives out of which the interpolating contour can be built (e.g., arcs of circles; Ullman, 1976); (ii) to specify an energy functional on the space of all possible interpolating curves, and seek the curve that minimizes that functional (e.g., Barrow & Tenenbaum, 1981; Horn, 1983; Mumford, 1994; Singh & Hoffman, 1999; Kimia et al., 2003) and (iii) to propose probabilistic models that combine stochastic generative models of contours, along with noisy estimates of geometric properties in the image to specify the probability for any interpolating curve (e.g., Feldman, 2001; Feldman & Singh, 2005; Singh & Fulvio, 2005; 2007; Williams & Jacobs, 1997; Yuille et al., 2004).

2. Psychophysical Studies

2a. Studies involving dot-sampled contours.

Warren et al. (2002, 2004) carried out studies investigating interpolation across gaps in 3D contours sampled at dots, viewed binocularly. Observers adjusted the position of an intervening probe dot so that it appeared to lie on the smooth contour that was perceived to pass through the dots. Warren et al. (2002) found that observers accurately interpolated both lines and parabolas with as few as 4-5 sample points, suggesting that these parametric contours were part of the "human visual spline". Other analyses examined how small perturbations of dot position affected contour interpolation (Hon et al., 1997; Warren et al., 2004). These revealed effects of only the two nearest dots on either side of the probe dot; interpolation settings were largely impervious to perturbations of farther dots.

These results suggest that the visual system combines information for interpolation relatively locally, incorporating that which is most necessary and reliable (Warren et al., 2002). This property of locality in interpolation is consistent with earlier results in contour grouping indicating that the visual system analyzes dot-sampled contours through local windows of four dots each (Feldman, 1997; 2001). Moreover, Warren et al. (2004) demonstrated that their results were best described by a model that minimizes the variance of the angles between neighboring line segments, which is consistent with smoothness effects observed in contour grouping (e.g., Pizlo et al., 1997).

2b. Extrapolation of partly-occluded contours.

Contour extrapolation is a critical component of the general problem of shape interpolation, since any interpolating contour must both smoothly extend (i.e., *extrapolate*) each of the two inducing fragments, as well as smoothly connect the two extrapolants (see, e.g., Ullman, 1976; Fantoni & Gerbino, 2003). In two studies, Singh & Fulvio (2005; 2007) examined contour extrapolation performance by mapping out the perceived trajectory of curved contours that disappeared behind an occluder, and analyzing the geometric characteristics of these visually extrapolated contours.

Observers viewed smooth contours that disappeared behind the straight edge of a half-disk occluder (see Figure 2a). They adjusted the position as well as the orientation of a short line probe on the opposite, curved side of the occluder, in order to optimize the percept of smooth continuation. For each inducing contour, extrapolation

measurements were obtained at six different distances from the point of occlusion by using half disks of different radii. Inducing contours used included linear segments, arcs of circles, parabolas, and Euler spirals (with linearly increasing, or decreasing, curvature; see Figure 2b).

Figure 2 about here

By fitting various shape models to the extrapolation data, and comparing these fits to the geometry of the inducing contours, Singh & Fulvio (2005; 2007) sought to characterize the geometry of visually extrapolated contours and its dependence on inducer geometry. The primary findings from these studies were that: (1) observers made systematic use of contour curvature in extrapolating contour shape; this has theoretical significance because current models of interpolation do not take into account inducer curvature; (2) visually extrapolated contours were characterized by decaying curvature with increasing distance from the point of occlusion; (3) the precision of observers' extrapolated contours decreased systematically with increasing curvature of the inducing contour; and (4) observers did not use rate of change of curvature in extrapolating smooth contours—despite the fact that they could detect it (as confirmed in a control experiment; Singh & Fulvio, 2007).

These properties served to characterize in formal terms the "good continuation" of smooth contours behind occluders. They were captured in a Bayesian model of contour extrapolation involving an interaction between two probabilistically defined constraints: a cocircularity-like tendency to continue the curvature of the inducing

contour, plus a prior for straightness (see the section on Models below).

2c. Interpolation of partly-occluded contours.

Fulvio, Singh, & Maloney (2008, 2009) extended the contour extrapolation paradigm to study visually interpolated contours. Part of the motivation was to understand how the visual extrapolants of two inducing contours are combined to produce a single interpolating contour. Fulvio et al. (2008) used a display in which the middle portion of a contour was occluded by a rectangular surface (see Figure 3). On each trial, a vertical "interpolation window" was opened at one of six possible locations, through which a short linear probe was visible. Observers adjusted the location (height) as well as orientation of the probe in order to optimize the percept of smooth continuation of a single contour behind the occluder.

Figure 3

The perceived interpolated contours were thus mapped out by taking measurements at six evenly spaced locations along the width of the occlusion region. The geometry of the inducing contours was varied—specifically, their turning angle (the angle through which one inducer must turn to align with the other) and their relative vertical offset. These manipulations affect the geometric relationship between the two extrapolants, and impose different boundary conditions on the interpolated contour. Geometric relations between the two inducers that are unfavorable—i.e., that violate the visual system's expectations about contours (see, e.g., Kellman & Shipley, 1991; Field,

Hayes, & Hess, 1993; Geisler et al., 2001)—should lead to failures of interpolation, i.e., an inability to produce a single, stable, and smooth interpolating contour.

Fulvio et al. (2008) established two behavioral measures of single, stable, smooth interpolation: (i) setting precision within each interpolation window; and (ii) internal consistency of the position settings and orientation settings across the multiple interpolation windows. They defined precision as the reciprocal of the standard deviation of observers' settings. Stable visual completion should allow the observer to repeatedly localize the contour over many trials, thereby yielding high setting precision. They defined internal consistency in terms of the mutual agreement between observers' positional settings and their orientation settings at the six measurement locations—in other words, the extent to which the combined set of measurements is consistent with a single, smooth contour. They measured this agreement in two ways: (i) parametrically, by comparing the observers' orientation settings to the tangent field of the best-fitting polynomial to the positional settings; and (ii) non-parametrically, by comparing the consecutive height differences implied by their orientation settings to the actual height differences observed between settings at consecutive locations.

The results showed that the turning angle between the two inducing contours primarily affected the precision of interpolation settings, with larger turning angles eliciting reduced precision. By contrast, the vertical offset between inducers adversely affected both precision and internal consistency. When the vertical offset was large enough that an inflecting curve would be required for smooth interpolation, they found that there is no single, stable and smooth contour that is consistent with observers'

interpolation settings of position and orientation at the multiple locations along the contour. This suggests a preference for smoother, lower-order polynomial interpolating contours, as well as the minimization of inflections (Takeichi et al., 1995; Singh & Hoffman, 1999).

Fulvio, Singh, & Maloney (2009) developed a purely experimental criterion for testing whether observers perceive a single, stable, and smooth interpolating contour for a given pair of inducing contours. Their "two-probe interpolation" paradigm avoids the use of explicit experimenter-defined criteria in characterizing performance in favor of relying solely on the observer's own performance. The structure of the experimental task was as follows: In Part I of the experiment, observers were asked to make repeated settings of interpolation position and orientation through multiple interpolation windows—one at a time, just as in Fulvio et al. (2008). In Part II, observers were asked to make the settings again, now with a second window opened, showing a fixed line probe which, unbeknown to them, was set to their own mean position and orientation setting from Part I of the experiment (see Figure 4a). If observers perceive a single, stable, smooth contour for any given inducer geometry, their performance should be identical in the two parts of the experiment. In other words, in Part II of the experiment, the presence of the fixed probe in the second window should not alter their interpolation settings relative to Part I—given that the fixed probe corresponds to their own setting of the interpolated contour.

Figure 4 about here

The results corroborated the previous finding in Fulvio et al. (2008) that inducing contour pairs requiring an inflecting interpolating contour do not yield the percept of a single, stable smooth contour. Under such conditions, observers' settings in Part II of the experiment systematically deviate from those in Part I, with a shift towards the nearest inducing contour—as if observers were instead interpolating between the fixed probe visible through the second window and the nearest inducing contour. Figure 4B shows such an example, with an observer's mean Part I settings depicted in red, and Part II settings depicted in blue. As can be seen, the Part II settings shift closer to the height of the nearest inducing contour. This influence of the nearest edges is consistent the earlier work described above (Feldman, 1997; Hon et al., 1997; Warren et al., 2004), which showed that integration and interpolation of dot-sampled contours are mostly influenced by local sources of information, as well as with Singh & Fulvio's (2005, 2007) finding of decrease in influence of an inducer's curvature with increasing distance from the point of occlusion.

The results of the two-probe task also revealed an improvement in setting precision (primarily for the positional settings) in Part II of the experiment, when the inducer geometry supported single, stable, smooth interpolation. This makes sense because the fixed probe segment acts as an additional consistent cue to interpolation. Counterintuitively, however, setting precision also improved in conditions that decidedly do not support contour interpolation, that is, for inducer geometries requiring inflecting interpolating contours. In this case, the improvement in precision most likely resulted from the shift in strategy noted above: observers switch to simply interpolating the

nearest visible segments, which makes the settings more precise but at the cost of sacrificing internal consistency.

The results of these two studies clarify the distinct roles played by the different components of the geometry of a pair of inducing contours. Although *turning angle* and *relative offset* between inducers are typically used jointly in defining contour "relatability" (see, e.g., Kellman & Shipley, 1991), our results show that these two components have qualitatively different effects on interpolation performance. The results also provide further evidence for interpolation being a selectively local process. In fitting polynomials to observers' interpolation settings, we found that conditions that produced failures in stable interpolation required polynomial fits of higher order than those that elicited successful, stable interpolation. Thus, smoother, lower-order, non-inflecting contours are clearly preferred by the visual system.

2d. The role of surface geometry in shape interpolation.

The studies reviewed so far investigated the role of contour geometry in shape interpolation. There is growing evidence, however, that shape interpolation is informed not only by contour geometry, but by surface (or region-based) geometry as well. Fulvio & Singh (2006) investigated the role of surface geometry on the perceived shapes of illusory contours, by embedding the same pairs of inducing contours in surfaces with different region-based geometries. Two types of region-based factors were manipulated: (i) *convex vs. concave:* a curved contour segment may correspond either to a convex protuberance on the surface, or to a concave indentation ("diamond" vs. "bowtie"; see

Figure 5a); and (ii) *width variation vs. axis curvature:* curvature along the bounding contour of a shape may arise either due to changes in the width of the shape across a straight axis, or from curvature on the medial axis of a shape with constant width (e.g., "diamond" vs. "bent tube"; see Figure 5b).

Figure 5 about here

Observers viewed the stimuli stereoscopically with the surfaces partially camouflaged by a rectangular surface (see Figure 5c). To the right of the illusory-contour display was an adjustable comparison surface that contained the same inducing contours connected smoothly by an adjustable middle portion. Observers adjusted the shape (degree of smoothness) of the contours of the middle portion so that they matched the perceived shape of the illusory contours in the interpolation display.

Although the local geometry of the inducing contours was equated, interpolation settings varied with changes to the region-based geometry: inducers embedded in convex ("diamond") surfaces were interpolated with flatter, smoother, curves than corresponding inducers embedded in concave ("bowtie") surfaces, which were interpolated with more peaked curves (i.e., closer to the point of intersection of the two linear extrapolants—one for each inducer). The influence of surface convexity corroborated earlier work on amodal completion of surfaces (Fantoni, Bertamini & Gerbino, 2005). In addition, however, Fulvio & Singh (2006) also found that the role of local surface convexity is systematically modulated by two other region-based geometric factors: (i) overall shape width: the effect of convexity is much more pronounced for

narrower shapes, and weakens considerably for wider shapes; and (ii) skeletal axis geometry: the effect of convexity is greater for “diamond” vs. “bow-tie” (straight axis with local modulations in width) and weaker for the convex vs. concave sides a “bent tube” shape (curved axis with constant width). These results demonstrate that ultimately contour interpolation must be understood in conjunction with the region-based geometry of the surface to which the contour belongs (see also Kogo et el., 2010).

3. Models

Approaches to modeling shape interpolation may be divided broadly into three categories:

1. Fixed class of shape primitives:

This category of models specifies a fixed class of shape primitives out of which the interpolating contour may be built. Ullman (1976), for example, modeled illusory contour shape using a pair of circular arcs that meet each other smoothly—with continuous tangents—and with each arc smoothly extrapolating an inducing contour. Out of the family of possible bi-arc solutions, the one with minimal total curvature is selected. (As we noted earlier, a special role for circular arcs has also been suggested by contour integration studies; e.g., Feldman, 1997; Pizlo, 1997). In addition to circular arcs, there is also evidence for the special status of *parabolas* in contour interpolation. Indeed, based on observers’ performance in interpolating dot-sampled contours, Warren et al. (2002) postulated that parabolas are part of the “human visual spline.”

2. Regularization or Variational Approaches:

This class of models specifies an energy functional on the space of all possible interpolating curves, and develops techniques to find the curve that minimizes this functional (e.g., Barrow & Tenenbaum, 1981; Horn, 1983). Two commonly used energy functionals are total curvature $\int \kappa(s)^2 ds$ (or bending energy), and total variation in curvature $\int \left(\frac{d\kappa}{ds} \right)^2 ds$. Minimizing total curvature along the length of the contour leads to a class of curves known as *elastica* (Mumford, 1994). By definition, these curves are "as straight as possible" given the boundary conditions of the interpolation problem, and the constraint of smoothness. Minimizing variation in curvature, on the other hand, penalizes *changes* in the curvature of the contour (rather than curvature itself). It is consistent with minimization of the variance in turning angles in the case of discretely sampled contours (e.g., Pizlo, 1997; Warren et al., 2004). It leads to a class of curves known as Euler spirals (Kimia et al., 2003), which are by definition as close to circular arcs as possible given the boundary conditions of the interpolation problem.

A more recent variational model (Ben-Yosef and Ben-Shahar, 2012) adopts a somewhat different strategy. Inspired by the columnar architecture of the primary visual cortex in primates, it takes as its underlying space $\mathbb{R}^2 \times S^1$—i.e., the space of all position-cum-orientations (x, y, θ). Given two inducing edges defined by (x_1, y_1, θ_1) and (x_2, y_2, θ_2)—i.e., the positions and orientations of the two inducers at their respective points of occlusion—the model seeks the shortest *admissible* curve in (x, y, θ)-space

that connects these two points in $\mathbb{R}^2 \times S^1$. A curve $\{(x(s),y(s),\theta(s)),s \in [0,L]\}$ in $\mathbb{R}^2 \times S^1$ is *admissible* if it defines an internally consistent smooth curve in the plane $\mathbb{R}^2$, i.e., if the orientations $\{\theta(s),s \in [0,L]\}$ are consistent with the tangent field of the curve $\{(x(s),y(s)),s \in [0,L]\}$ in $\mathbb{R}^2$. The authors argue that this shortest-length constraint on admissible curves in $\mathbb{R}^2 \times S^1$ may be interpreted in terms of minimal neural activation in the primary visual cortex—hence the "minimal energy" constraint in this case is meant to be biologically motivated.

Although not a variational model in the technical sense, Fantoni & Gerbino's (2003) model is nevertheless variational in spirit. In their model, the shape of the interpolating contour results from the interaction between two constraints: (i) a "good-continuation" (GC) constraint for the interpolating contour to smoothly continue—i.e., extrapolate—each respective inducer; and (ii) a "minimal-path" (MP) constraint for the interpolating contour to be as short as possible. Each constraint is used to define a vector field in the intervening space between the two inducers. The GC constraint by itself would extend each inducer smoothly; however this generically results in a tangent discontinuity at the locus of intersection of the two extrapolants. The MP constraint by itself would simply connect the two points of occlusion with a straight line segment—generally resulting in a tangent discontinuity at each point of occlusion. The interaction of these two constraints generates a smooth interpolating contour that lies somewhere between these two extremes.

3. Probabilistic models:

Probabilistic models combine generative models of contours, along with noisy estimates of geometric properties in the image to specify the probability for any interpolating curve. Generative models of contours are often expressed as probabilistic distributions on turning angles. That is, given a contour specified up to a certain point, where is it likely to go next? How much is it likely to "turn" from its current tangent direction? Psychophysical evidence suggests that, for open contours, the visual system implicitly assumes that the contour is most likely to go straight; i.e., a turning angle of 0 is most likely, with probability falling off with the magnitude of the turning angle (e.g., Field et al., 1993; Feldman, 2001; Geisler et al., 2001; Yuille et al., 2004). This is often modeled as a von Mises distribution centered on 0 (e.g., Feldman & Singh, 2005; Singh & Feldman, 2012). A model of contour interpolation based on this generative model would on average yield solutions similar to those of variational models based on minimizing curvature (e.g., Williams & Jacobs, 1995).

The results of contour extrapolation studies suggest a slightly different approach, however. Recall that in these studies, observers made extrapolation settings not only in the immediate vicinity of the point of occlusion, but at multiple distances away from it (Singh & Fulvio, 2005; 2007). One of their main findings was that visually extrapolated contours were consistently characterized by decaying curvature with increasing distance from the point of occlusion. This behavior was modeled as a Bayesian interaction between two constraints—both expressed as probability distributions on extrapolation curvature. The prior term captures the default expectation of the visual system that at

any given point a contour is most likely to go "straight" (this is the same constraint used in generative models of contours considered above). It is modeled as a Gaussian distribution on curvature that peaks at 0, with some fixed variance. The likelihood term captures the tendency of visually extrapolated contours to continue the estimated curvature of the inducing contour at the point of curvature. Thus, given an inducing contour, the likelihood is a distribution centered on the estimated curvature of the inducing contour. However, its spread is not fixed, but rather increases monotonically with distance from the point of occlusion.[1] Under the assumption of Gaussian distributions, the MAP estimate of extrapolation curvature is a weighted average of: (i) the estimated curvature of the inducing contour at the point of occlusion, and (ii) zero curvature (the expected value of the prior distribution). Importantly, the weights assigned to these two components are inversely proportional to their respective variances. Near the point of occlusion, the likelihood (i.e., the continuation of estimated inducer curvature) is subject to very little noise, and thus the likelihood dominates the prior. With increasing distance from the point of occlusion, however, the reliability of the likelihood decreases monotonically (as a result of increasing spread), and thus the prior—which has a fixed variance—eventually comes to dominate the likelihood. Since the prior is centered on 0 curvature, the weighted averaging predicts a systematic decrease in extrapolation curvature with increasing distance from the point of occlusion—namely, the decaying curvature behavior.

[1] Singh & Fulvio (2005) assumed a linear increase in standard deviation with increasing distance from the point of occlusion, consistent with a Weber-like behavior. The decaying curvature behavior, however,

There is an important sense in which this probabilistic model of contour extrapolation differs from regularization/variational models considered earlier as well as stochastic generative models of contours. In both of these classes of models, the same constraint (e.g., minimization of curvature) is assumed to apply uniformly along the entire length of the contour. Whereas this analysis of visually extrapolated contours suggests that human vision adopts different shape constraints—at least in the sense of giving different weights to different shape constraints—depending on distance from the point of occlusion. For instance, it gives greater weight to minimization of variation in curvature (or the continuation of inducer curvature) near the point of occlusion, but gives greater weight to minimizing curvature at greater distances from the point of occlusion.

If this is true more generally, it suggests that variational models—which implicitly assume that a single shape constraint applies along the entire length of the interpolated contour—are unlikely to capture the shape of visually interpolated contours. Rather, the shape constraints used (or the weights assigned to them) are likely to vary with distance from the nearest point of occlusion. It will be of interest to develop a new probabilistic model of visual contour interpolation within this framework, and test its predictions against the large amount of data on contour interpolation summarized here.

---

requires only an assumption of a monotonic increase.

References

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

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

Fantoni, C., Bertamini, M., & Gerbino, W. (2005). Contour curvature polarity and surface interpolation. *Vision Research, 45*, 1047–1062.

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

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

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

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

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

Fulvio, J.M., Singh, M., & Maloney, L.T. (2008). Precision and consistency of contour interpolation. *Vision Research, 48,* 831–849.

Fulvio, J.M., Singh, M. ,& Maloney, L.T. (2009). An experimental criterion for consistency in interpolation of partially-occluded contours. *Journal of Vision, 9(4):5*, 1-19.

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

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

Grossberg, S., & Mingolla, E. (1985). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review, 92*(2), 173–211.

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.

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

Hume, D. (1748/1993). *An enquiry concerning human understanding*. Indianapolis, IN: Hackett Publishing.

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

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

Kogo, N., Strecha, C., Van Gool, L., Wagemans, J. (2010). Surface construction by a 2-D differentiation-integration process: A neurocomputational model for perceived border ownership, depth, and lightness in Kanizsa figures. *Psychological review, 117 (2),* 406-439.

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

Parent, P. & Zucker, S.W. (1989). Trace inference, curvature consistency, and curve detection. *IEEE Transactions on pattern analysis and machine intelligence, 2*(8), 823-839.

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

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

Singh, M., & Fulvio, J.M. (2007). Bayesian contour extrapolation: Geometric determinants of good continuation. *Vision Research, 47,* 783–798.

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

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

Takeichi, H., Nakazawa, H., Murakami, I., & Shimojo, S. (1995). The theory of the

curvature-constraint line for amodal completion. *Perception, 24,* 373–389.

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

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

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

Wertheimer, M. (1923), Untersuchungen zur Lehre von der Gestalt II. *Psychologische Forschung*, *4*, 301-350. Translation published in Ellis, W. (1938). *A source book of Gestalt psychology.* London: Routledge & Kegan Paul (pp. 71-88).

Williams, L. R., & Jacobs, D. W. (1997). Stochastic completion fields: A neural model of illusory contour shape and salience. *Neural Computation, 9,* 837–859.

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

Figure Captions

Figure 1. (a) Illustration of the *grouping or segmentation problem*. Here, the visual system must determine which of the four contours on the right side of the occluder continues the contour on the left side. (b) Illustration of the *shape problem*. The visual system must decide how to complete the connection between two grouped contours.

Figure 2. (a) Schematic of the stimulus displays used by Singh & Fulvio (2005; 2007). A contour segment ("inducer") is occluded by the midpoint of the straight edge of a half disk occluder, which can be one of six radius sizes. A short line segment ("adjustable probe") resides at the opposite, curved edge of the occluder. Observers adjust both the position and orientation of the probe so that it appears to smoothly extrapolate the inducer. (b) Types of contours used in the studies: linear, parabolic, circular, and Euler spiral of increasing and decreasing curvature, respectively.

Figure 3. Schematic of the stimulus display used by Fulvio, Singh, & Maloney (2008). Two linear inducers are occluded by opposite edges of a rectangular occluder. An adjustable probe is visible through a small slit ("interpolation window") in the occluder. Observers adjust the position and orientation of the probe so that it appears to smoothly interpolate the two inducers.

Figure 4. (a) Schematic of the stimulus display used by Fulvio, Singh, & Maloney (2009). Two linear inducers are occluded by opposite edges of a rectangular occluder. An additional inducer, set to the observer's mean position and orientation settings in the single window version of the task (recall Figure 3), is visible through an interpolation window. An adjustable probe is visible through a second interpolation window.

Observers once again adjust the position and orientation of the probe so that it appears to smoothly interpolate the two flanking inducers. (b) Representative data depicting a shift in the mean probe settings towards the nearer flanking inducers when an inducer in a second interpolation is present.

Figure 5. (a) Schematic of the surface stimuli used in Fulvio & Singh (2005). The contours defining the surfaces are identical, however, their concavity varies. Top: convex diamond; Bottom: concave bowtie. (b) Additional stimulus schematic. Once again, the contours defining the surfaces are identical, but their axis geometry varies. Top: convex diamond; Bottom: bent tube. (c) Depiction of the 3D experimental stimuli from the observer's standpoint.

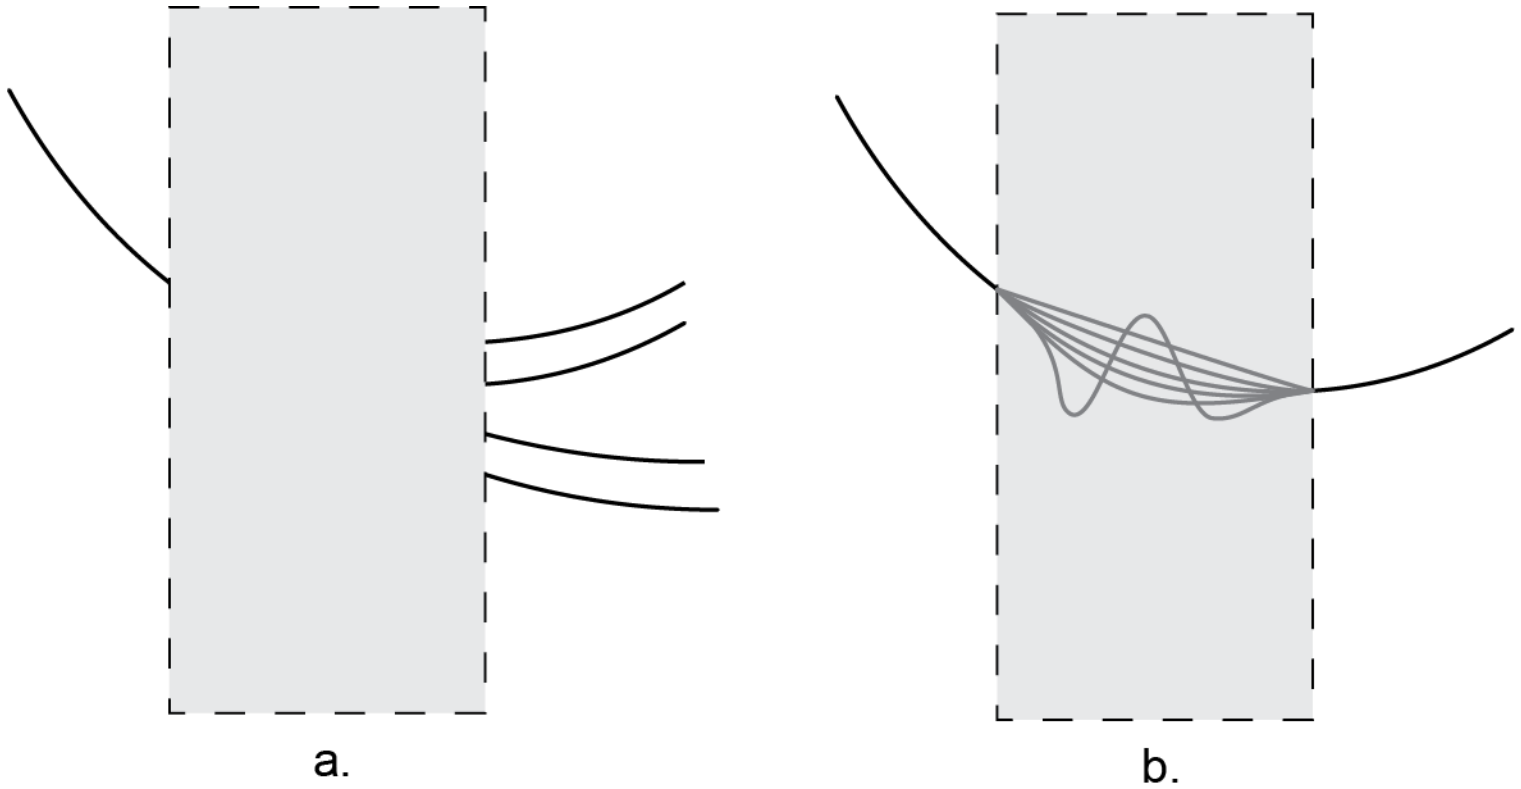

Figure 1.

a.

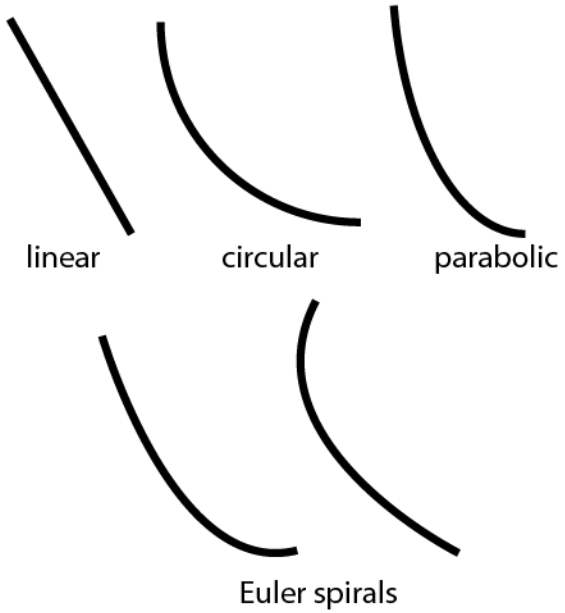

b.

Figure 2.

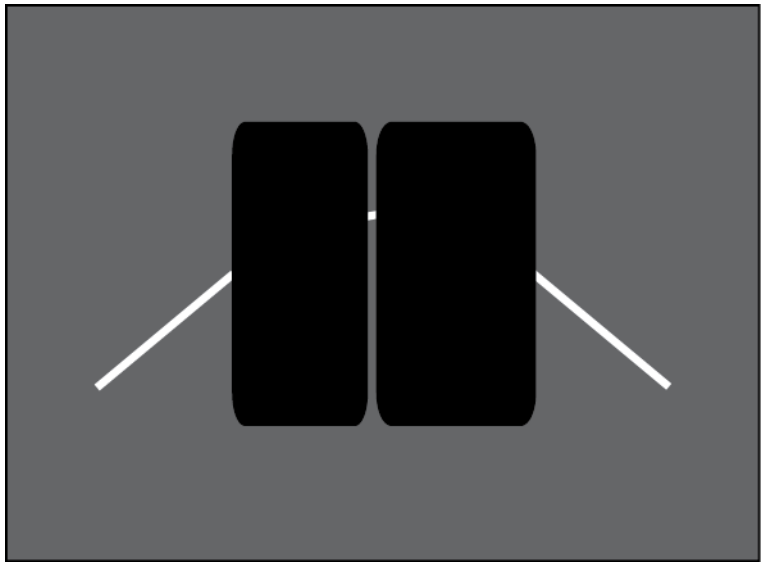

Figure 3.

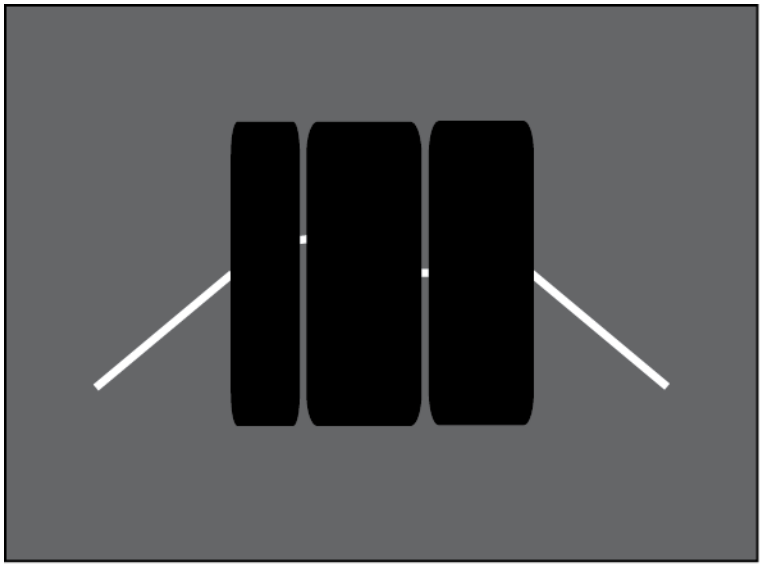

a.

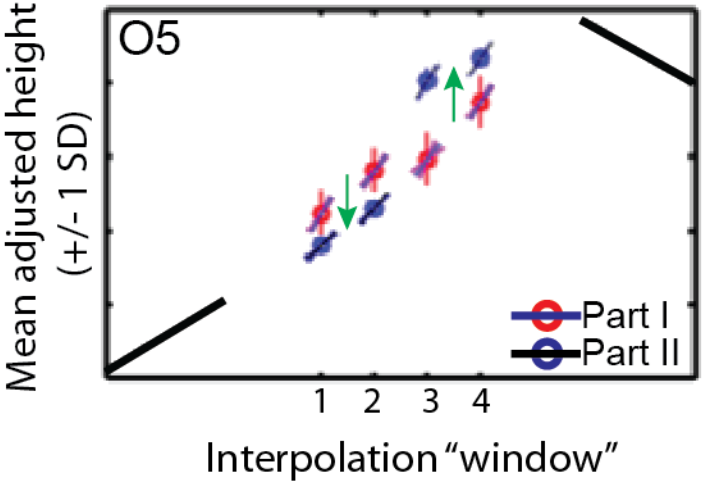
O5
Mean adjusted height
(+/- 1 SD)
Part I
Part II
1
2
3
4
Interpolation "window"

b.

Figure 4

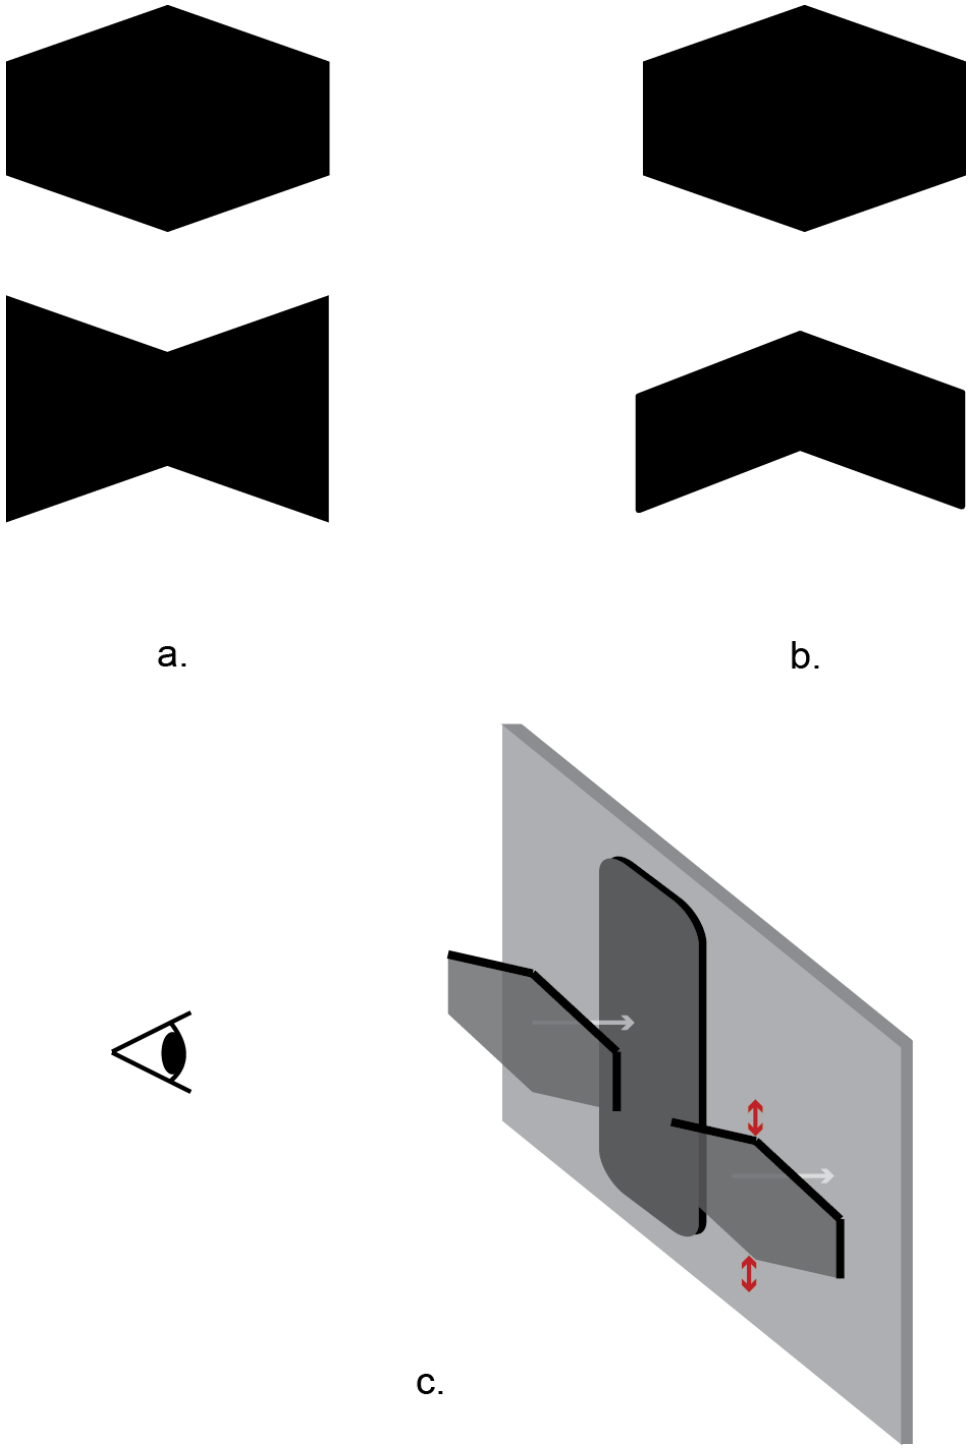

Figure 5.