



A unified model of illusory and occluded contour interpolation

Donald J. Kalar^{a,*}, Patrick Garrigan^{b,1}, Thomas D. Wickens^c, James D. Hilger^a, Philip J. Kellman^{a,2}

^a Department of Psychology, University of California, Los Angeles, CA, United States

^b Department of Psychology, Saint Joseph's University, Philadelphia, PA, United States

^c Department of Psychology, University of California, Berkeley, CA, United States

ARTICLE INFO

Article history:

Received 21 April 2009

Received in revised form 8 October 2009

Keywords:

Illusory contours
Occluded contours
Identity hypothesis
Contour interpolation
Computational modeling

ABSTRACT

Models of contour interpolation have been proposed for illusory contour interpolation but seldom for interpolation of occluded contours. The identity hypothesis (Kellman & Loukides, 1987; Kellman & Shipley, 1991) posits that an early interpolation mechanism is shared by interpolated contours that are ultimately perceived as either illusory or occluded. Here we propose a model of such a unified interpolation mechanism for illusory and occluded contours, building on the framework established in Heitger, von der Heydt, Peterhans, Rosenthaler, and Kubler (1998). We show that a single, neurally plausible mechanism that is consistent with the identity hypothesis also generates contour interpolations in agreement with perception for cases of transparency, self-splitting objects, interpolation with mixed boundary assignment, and “quasimodal” interpolations. Limiting cases for this local, feed-forward approach are presented, demonstrating that both early, local interpolation mechanisms and non-local scene constraints are necessary for describing the perception of interpolated contours.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

A basic problem in seeing is occlusion. In ordinary scenes, only some parts of an object reflect light that reaches the eyes, and very often, parts of an object are not visible because another object or surface is positioned between the object and the observer. Moreover, a given object may project to the eyes in multiple, separated regions, as when we view objects through foliage. Despite this fragmentary input, typical visual experience is of complete objects. For unified objects to be perceived under partial occlusion, visual processes must connect the spatially separated parts of the scene that correspond to the same physical objects.

Evidence suggests at least two distinct processes operate to accomplish object formation. Perhaps the most fundamental process is contour interpolation, which refers to the representation of contour segments not locally specified in the optical input. The other process is known as surface interpolation; under occlusion, surface properties propagate behind occluders to link visible regions (Yin, Kellman, & Shipley, 1997). These processes are complementary in that surface interpolation appears to be confined by real or interpolated contours, when these are present (Kellman, 2003; Yin et al., 1997), although surface interpolation can also operate in their absence. The visual system's representation of

the occluded portion of objects includes both shape and surface properties that are determined by the shape and surface properties of the visible portions of objects.

Researchers have made considerable progress in understanding a number of aspects of interpolation processes (see below). In this paper, we focus on physiologically-inspired modeling of contour interpolation. Research has led to an understanding of the conditions under which interpolation occurs, but we know less about how the geometric properties of contour interpolation are realized in neural mechanisms. Much of what we do know comes from experiments on illusory and occluded contour formation. A combination of psychophysical data, theoretical considerations, and recent physiological evidence suggest that a common interpolation mechanism participates in otherwise different phenomena, specifically illusory contours and completion of partly occluded contours (Albert, 2007; Kellman, Garrigan, & Shipley, 2005; Kellman, Yin, & Shipley, 1998; Murray, Foxe, Javitt, & Foxe, 2004; Ringach & Shapley, 1996).

Illusory contours (ICs) can occur when the visual properties of an occluding surface are similar to the background in the scene and the visible portions of the occluded surface's bounding contours meet a specific set of geometric constraints (see Section 2.1, below). Under these conditions, observers report perceiving continuous contours bounding the occluding object, despite the absence of any physically specified bounding contour along parts of its perimeter.

Current models of contour interpolation vary in whether they posit a shared interpolation mechanism for illusory and occluded

* Corresponding author.

E-mail address: donkalar@ucla.edu (D.J. Kalar).

¹ These authors contributed equally.

² This research was supported by NIH EY13518-01 to PJK.

contours. Some are consistent with this notion (e.g., Grossberg, 1994), whereas others are explicitly aimed at either illusory (e.g., Heitger, von der Heydt, Peterhans, Rosenthaler, & Kubler, 1998; Ullman, 1976) or occluded contours (e.g., Buffart, Leeuwenberg, & Restle, 1981). The model presented in this paper is adapted from a well-specified and plausible feed-forward, neurally-inspired model of IC interpolation proposed by Heitger et al. (1998). This model addresses formation of illusory, but not occluded, contours. This limitation was intentional, motivated by neurophysiological data suggesting that some cells in monkey area V2 play a role in IC formation, but did not seem to be engaged by a corresponding occlusion display. Subsequent physiological and psychophysical work has produced mixed results regarding the neural locus of illusory and occluded contour formation (e.g., Bakin, Nakayama, & Gilbert, 2000; Hirsch et al., 1995; Kellman, Garrigan, & Shipley, 2005; Lerner, Hendler, & Malach, 2002; Mendola, Dale, Fishl, Liu, & Tootell, 1999; Sugita, 1999).

Here we develop a model of contour interpolation that completes both illusory and occluded contours. The model does not center on the question of where in the visual system contour interpolation occurs, but rather focuses on the nature of neural coding that first represents visual scenes and the initial operations of interpolation that allow contour connections to surmount gaps in the input. Building on the work of Heitger and colleagues, we have developed an operator that detects the local features that indicate occlusion and interpolates contours through the occluded regions. Compared to the model put forth by Heitger and colleagues, this model is more general, and is able to match human perception for a larger set of occlusion stimuli. The primary advance in this model is that it handles both occluded and illusory contours. While it is designed for the purpose of completing partially occluded contours, it also generates ICs that are consistent with human perception. We believe the unified treatment of the basic interpolation stage in illusory and occluded contour formation results in a more general and accurate account. The new model, however, also helps to reveal the limitations of local, feed-forward, non-symbolic processing when contrasted with human phenomenology. We find that while our local, feed-forward model does generate ICs consistent with perception for many cases, it cannot account for several specific classes of IC displays. In these cases it can be shown that higher-level, non-local scene information is necessary to account for human perception. The simulations reveal some of the higher-level processing requirements that must extend beyond local, image-based contour computations.

2. Background

Researchers have applied a number of different approaches to understanding contour interpolation. One approach has been to describe the conditions under which interpolation does and does not occur (e.g., Kellman & Shipley, 1991; Kellman, Garrigan, Yin, Shipley, & Machado, 2005). Other research has focused on the shapes of interpolation paths, based on geometric constraints (e.g., Fantoni & Gerbino, 2003; Ullman, 1976), or derived by considering how interpolation might occur through a restricted set of operations (e.g., Heitger et al., 1998). These approaches complement each other, in that they tell us different things about interpolation processes. We briefly review research describing the geometric conditions that support contour interpolation and the relation between occluded and illusory contour interpolation.

2.1. Geometric conditions for contour interpolation

What visible areas are linked by the visual system to form representations of complete objects? The question is fundamental,

because adjacent areas in the optical projection often belong to different objects or surfaces, due to occlusion. Moreover, a single object may project to multiple, spatially separated visible regions. Both the initiating conditions and geometric relationships underlying object interpolation have become better understood in recent years. Object interpolation appears to depend on complementary processes of contour and surface interpolation (Grossberg & Mingolla, 1985; Kellman & Shipley, 1991).

The initiating conditions for contour interpolation are contour junctions in the optical projection, of which there are two primary types: T-junctions (often given by an edge passing behind an occluding surface) and L-junctions (sharp corners formed from two joining contours). Shipley and Kellman (1990) observed that in general, interpolated contours begin and end at these junctions in images, and showed that their removal eliminated or markedly reduced contour interpolation. It has been proposed that second-order discontinuities may be capable of weakly triggering interpolation (Albert, 2001; Albert & Hoffman, 2000; Shipley & Kellman, 1990). This may be the case, or such effects may arise from detection of tangent discontinuities at multiple scales (Albert, 2001; Wurtz & Lourens, 2000). A corner that is slightly rounded in a display may be detected as a tangent discontinuity by a coarse (low spatial frequency) detector. Heitger, Rosenthaler, von der Heydt, Peterhans, and Kubler (1992) labeled the location of contour junctions “key-points” and proposed a neurally plausible model for their extraction from images. Part of the challenge in using key-points to initiate contour interpolation is that contour junctions are necessary for contour interpolation but are not sufficient. Many corners in images are corners of objects, not points at which a contour passes behind an intervening surface (or in front, as in ICs).

Whether contour interpolation occurs between junctions depends on a set of geometric relations that have been formalized as the theory of *contour reliability* (Kellman & Shipley, 1991; Singh & Hoffman, 1999). According to contour reliability, the visual system interpolates when two visible contours, ending at junctions, can be connected by an interpolated contour that is smooth (differentiable at least once), monotonic, and bends through no more than about 90°. Fig. 1 shows, for a single edge, the range of orientations at a particular position within which contour interpolation will occur. In the example shown, an edge defined in 2-D Cartesian coordinates is oriented along the x axis and ends at the origin. For this edge, the set of reliable orientations (θ , defined as $\theta = 0$ for edges oriented horizontally) at positions $(x \geq 0, y \geq 0)$ is:

$$\tan^{-1} \frac{y}{x} \leq \theta \leq \frac{\pi}{2}$$

and for $(x \geq 0, y \leq 0)$ is:

$$\tan^{-1} \frac{y}{x} \geq \theta \geq \frac{\pi}{2}$$

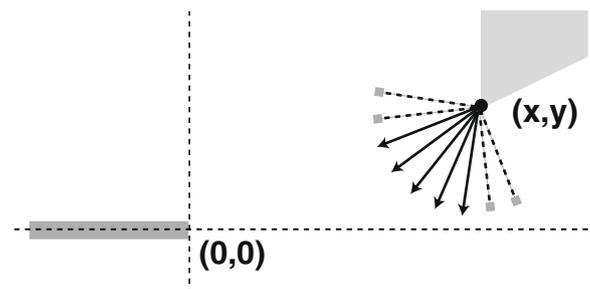


Fig. 1. Reliability in 2-D Cartesian coordinates. For an edge oriented along the x axis and ending at the origin (bottom, left), the set of reliable orientations for an edge ending at point (x,y) all lie within the gray shaded area (top, right). Pointed arrows indicate the orientation of potential interpolation paths at point (x,y) .

This relation defines the limits of relatability, i.e., a range of positions and orientations of two edges within which contour interpolation may cause them to be perceived as connected. The importance of defining limits is that contour connections in object perception often involve discrete categorization: Visual processes determine whether two visible fragments are or are not parts of the same object. The classification of visible parts as being connected or not is probably decisive in determining whether certain kinds of further processing occur (e.g., generation of a representation of connected areas, generation of a single shape description encompassing those parts, comparison with items or categories in memory). Recent research indicates that the representation of visual areas as parts of a single object or different objects has a variety of important effects on information processing (Baylis & Driver, 1993; Zemel, Behrmann, Mozer, & Bavelier, 2002).

2.1.1. Ecological foundations

Relatability is related to, but distinct from the Gestalt principle of good continuation (Wertheimer, 1923). Both embody underlying assumptions about contour smoothness, but they take different inputs and have different constraints (Kellman, Garrigan, Kalar, & Shipley, 2003). Good continuation in Wertheimer's demonstrations applied to the breakup of contiguous contours, such as three intersecting line segments. Relatability applies to spatially separated segments, and neither the requirements for smoothness across gaps, monotonicity, nor a limit around 90° of turn angle apply to cases of Wertheimer good continuation. In fact, the contour segments leading into junctions in Wertheimer displays may be non-monotonic and may bend through any angle while maintaining good continuation. The common nexus of Wertheimer good continuation and relatability is the first-order contour discontinuity (sharp corner). Although this point has seldom been made formally, the presence of a first-order or tangent discontinuity appears to govern the breakup of Wertheimer displays (as in determining which ways two road lines on a map continue when they intersect). Whereas a first-order discontinuity breaks Wertheimer continuation, this image feature is in fact the trigger for the operation of relatability, which applies to connecting separated edge fragments across gaps. Thus, these related principles have in common a notion of smoothness that hinges on first-order discontinuities; their presence in an image signals the end of contiguous image fragments but may also signal locations where interpolation across gaps can occur.

Image statistics indicate that the geometry of contour relatability may capture an optimal approach for contour interpolation in natural scenes. Geisler, Perry, Super, and Gallogly (2001) found that the spatial distribution of pairs of edge elements from a single contour largely agreed with the geometry of relatability. Two visible edge segments from the same physical contour meet the mathematical relatability criterion far more often than not. This geometry also shows up in behavioral measures of contour grouping, where subjects detect paths formed from Gabor elements concealed among similar elements with random orientation (Field, Hayes, & Hess, 1993). They propose that contours are linked by an "association field", formed from interactions among orientation-sensitive units in V1. Yen and Finkel (1998) proposed a model based on properties of similarly interacting cortical units for extraction of contours from images.

2.2. The identity hypothesis

Contour interpolation results in two distinct phenomena: ICs and occluded contours. These may be distinguished phenomenologically. Michotte, Thines, and Crabbe (1964) called illusory contours and surfaces "modal completion" because their appearance

included sensory attributes (modes). For a modally completed display, one could answer the question of whether the completed surface looks brighter than the background, for example. Michotte et al., called occluded contours generated by completion processes "amodal", in that they were clear perceptual descriptions but did not manifest sensory attributes. The two outcomes of interpolation may also be distinguished functionally (e.g., Kellman, 2003). "Modal completion" (illusory contours and surfaces) labels cases where interpolated contours and surfaces appear nearest to the observer (i.e., unobstructed by nearer surfaces). "Amodal completion" (occluded contours and surfaces) labels cases in which interpolated parts are represented as positioned behind some intervening surface. From this functional standpoint, it is natural to consider these interpolation cases in a unified fashion: The visual system interpolates objects across gaps, and these perceptually completed objects can be in front of or behind other objects, depending on information in the scene. The difference in appearance for "modal" and "amodal" simply marks whether the completed object is in front of or behind other surfaces. This idea of a unified interpolation process, allowing differently appearing outcomes, is consistent with findings that both interpolation appearances operate within the geometric constraints described above. Other similarities and logical arguments (presented below), have led to the idea that illusory and occluded contour interpolation depend on a common underlying mechanism during early stages in processing.

The formal similarities of illusory and occluded contours are illustrated in Fig. 2. At the top are shown an illusory and an occluded display with equivalent visible edges and gaps, leading to the perception of similarly shaped completed objects. Fig. 2, bottom, shows two self-splitting figures in which a homogeneous area is seen as splitting into two objects. In each of these cases, visual processes connect areas across gaps, and appear to confer specific contour shapes in the unspecified regions.

The claim that, despite their phenomenological differences, these completion phenomena may share a common underlying interpolation mechanism has been called the *identity hypothesis* (Kellman & Loukides, 1987; Kellman & Shipley, 1991). This claim

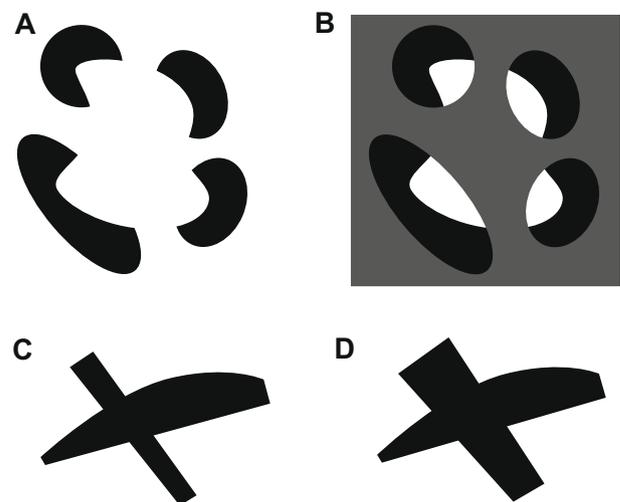


Fig. 2. Contour completion. Equivalent displays that generate illusory (A) and occluded (B) contours are shown. The completed surfaces, after contour interpolation, appear to have the same shape. Self splitting figures (bottom) also involve contour interpolation. The perceived depth order of two surfaces is determined by the length of crossing interpolation paths, with the shorter path causing the associated surface to appear in front (with ICs), and the longer path causing the associated surface to appear behind (with occluded contours). This is evident in (C). In (D), the interpolations are approximately equal in length, and the perception of depth order may appear bistable.

adopts the view of Michotte and colleagues that both modal and amodal completion are perceptual phenomena, and suggests that a common interpolation processes connects fragments across gaps in both. These connections take place at an early stage of processing, and help organize scenes into connected units. Relative to the observer, contour interpolation behind an intervening surface is amodal and the completed contour will appear occluded. Contour interpolation in front of an intervening surface is modal and the completed contour will appear illusory. A single interpolated contour or surface may even appear modal along some of its extent and amodal at other places (Kellman, Garrigan, & Shipley, 2005; Kellman, Garrigan, Shipley, & Keane, 2007). In terms of processing, the modal or amodal appearance results from depth information and scene constraints that determine depth order of the objects and surfaces in a scene.

Depth information can constrain interpolation at the outset, as shown in studies of 3D interpolation (Kellman, Garrigan, Yin, et al., 2005). Computations of depth relations among surfaces must also incorporate the outputs of interpolation (Fantoni & Gerbino, 2008; Kellman et al., 1998; Kellman, Garrigan, & Shipley, 2005). Phenomenologically, the difference in appearance of amodal and modal completions corresponds to their depth position in front of or behind other opaque objects and appears to mark this important attribute (e.g., it may determine whether the observer can reach unimpeded for an object). In this view, modal or amodal appearance has more to do with the positions of other surfaces relative to completed objects than with the underlying completion process per se.

This idea is consistent with the shifting of modal and amodal appearance in displays with little or no depth information, as in Fig. 2D (see discussion below). The identity hypothesis is also consistent with psychophysical evidence indicating that illusory and occluded contour interpolation have similar determinants (Gold, Murray, Bennett, & Sekuler, 2000; Kellman et al., 1998; Ringach & Shapley, 1996; Shipley & Kellman, 1992b), time course (Guttman & Kellman, 2004) and strength (Kellman et al., 1998; Ringach & Shapley, 1996), as well as with recent neurophysiological results (Murray et al., 2004). Other research has suggested differences in constraints on illusory and occluded contour interpolation (Anderson, 2007; Anderson, Singh, & Fleming, 2002) or their neural substrates (Corballis, Fendrich, Shapley, & Gazzaniga, 1999; von der Heydt, Peterhans, & Baumgartner, 1984). These issues have been discussed in detail elsewhere (Albert, 2007; Kellman, 2003; Kellman, Guttman, & Wickens, 2001; Kellman et al., 2007).

Strong evidence for the identity hypothesis has taken at least two forms. The first involves examples where a single interpolated

contour connects illusory and occluded portions. An example is shown in Fig. 3. Evidence from an objective performance paradigm shows that these interpolated contours have effects indistinguishable from ordinary illusory and occluded contour completion (Kellman et al., 1998). Typical ICs require inducing elements on both sides of a gap (e.g., Grossberg & Mingolla, 1985; Heitger et al., 1998), as do completed occluded contours. None of the interpolations in Fig. 3 fulfill the requirements for either illusory or occluded contour completion, because each has an inducer for an occluded contour at one end and an IC at the other. Interpolation nevertheless occurs, and each completed contour appears as illusory along part of its extent and occluded along the rest. These connections between illusory and occluded inducers, which Kellman et al. (1998) called *quasimodal* contours, naturally fit with the notion of a unified interpolation process. Outputs of interpolation can appear in front of or behind other surfaces, even for the same interpolated contour.

Strong support for the identity hypothesis also comes from cases in which interpolation appears to precede the determination of the illusory or occluded appearance. Such phenomena are consistent with the idea of an early interpolation process that can lead to differing perceptual appearances. The shapes in Fig. 2C and D are examples of a class of displays studied by Petter (1956). Although these shapes are formed from a contiguous, homogeneously colored and textured area, perceptually they appear as two objects. The splitting into two objects requires corners and relatable edges, and is predicted by the object formation model of Kellman and Shipley (1991). Note that despite the fact that all parts of this display lie in the same depth plane, the visual system appears to impose a constraint that one object must appear in front. In the areas where the two objects cross, neither has contours given by local image information; the perception of complete boundaries for each form depends on interpolated edges. Due to the depth ordering, the nearer display has ICs, whereas the farther display has occluded contours. Petter (1956) noted that where the interpolated boundaries cross, the boundary that traverses the smaller gap appears in front. According to Petter's rule, whether an interpolated contour appears in front or behind – i.e., as illusory or occluded – depends on its length relative to the interpolated contours that it crosses. Logically, this means that the determination of illusory or occluded appearance involves some comparison or competition of crossing interpolations. In such cases, interpolation (or more precisely, a relation involving a pair of interpolations) determines the illusory or occluded appearance of completed contours. The alternative to the identity hypothesis – that there are wholly separate processes of illusory and occluded contour formation (Anderson, 2007) – cannot explain the dependence of modal or

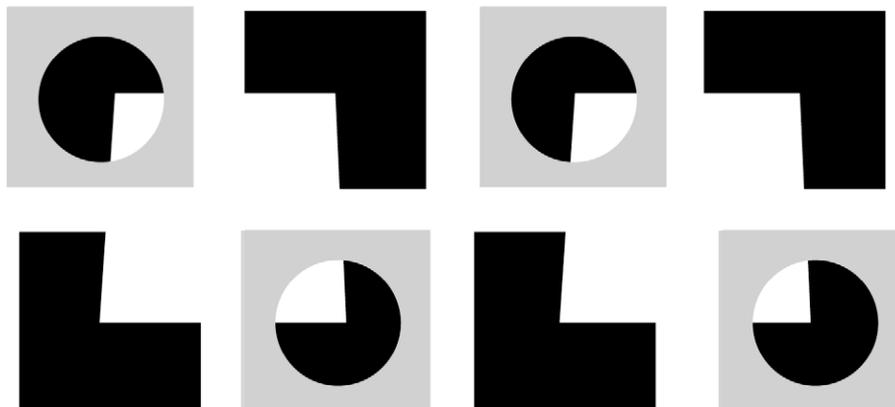


Fig. 3. Quasimodal interpolation. Display is a stereo pair, which may be free-fused by crossing the eyes. Interpolation connects IC inducing elements to occluded contour inducers (see text).



Fig. 4. Object formation, depth spreading, and modal/amodal appearance. When cross-fused, the rectangle appears to pass in front of the vertical bar on the left and behind the vertical bar on the right. The “rectangle” and the “vertical bars” have complete boundaries only because of interpolation; in fact, these are crossing interpolations. Where the rectangle passes behind, its contours are occluded and where it passes in front, it has ICs. The depth orderings and modal or amodal appearance of the ICs are consequences of interpolation. The reason is that the depth of each vertical bar relative to the rectangle depends on the depth values of the rectangle at the crossing points. The rectangle’s depth, however, must be inherited via depth spreading from its vertical edges (the only features available for extracting disparity). But depth spreading is confined within objects; it does not propagate unchecked through the entire display. Since the object’s boundaries have been formed by interpolation, and depth spreading depends on the object boundaries, both the perceived depth order and the modal/amodal appearance of several contours in the display are consequences of, not precursors to, interpolation processes.

amodal appearance on the relative extent of crossing interpolations (Kellman et al., 2007).

Petter effect displays are not the only phenomena in which interpolation logically precedes determination of illusory or occluded appearance. Fig. 4 displays a stereogram which, when fused, shows a central rectangle slanted in depth about a vertical axis with two occluding surfaces. One surface appears in back and partly occluded, the other appears in front and partly illusory. In this type of display illusory or occluded appearance is a consequence of depth spreading. Depth information spreads within objects (here because, it is given by stereoscopic information only at the left and right edges of the central white areas). Since the central white object is established in part by contour interpolation, interpolation must precede determination of illusory or occluded appearance (for more detailed discussion, see Kellman, Garrigan, & Shipley (2005)).

2.3. Computational approaches to interpolated contour shape

As discussed above, geometric information available in the image predicts when interpolation occurs, and consequently, researchers have proposed a number of “bottom-up” models of contour interpolation. Typically such models find an interpolation solution that minimizes some geometric quantity. Ullman (1976) proposed a model of this type in which the interpolation solution was composed of two circular arcs smoothly extending from their respective contours, with equal slopes at the point where they meet. The specific pair of arcs chosen was the one with the lowest total curvature, $\int (d\alpha/dl)^2$, where α is the slope of the interpolated contour with length l .

Other models have been proposed that do not have specific shape primitives, allowing more flexibility in the interpolations they generate. Williams and Jacobs (1997) proposed a model in which a unique interpolation solution between two edges is found by treating one edge as a source, the other as a sink. Many interpolation paths between these points are generated by moving from source to sink through a random-walk procedure, and the average random path is chosen. Fantoni and Gerbino (2003) proposed a model in which paths that minimize the distance between the endpoints of the occluded contours compete with paths that are linear (or curvilinear) extensions of the occluded contours. The

“minimum path” field and the “good continuation path” field combine to form smooth interpolation solutions for many occluded contour geometries.

2.4. Modeling neural mechanisms of contour interpolation

Other researchers have built models of contour interpolation assuming the outputs of neural units with well understood receptive fields as input. Heitger et al. (1992) presented a model in which images are convolved with simulated simple and complex cells, the outputs of which are combined to locate corners and contour junctions. Later, Heitger et al. (1998) showed that these outputs could be grouped, depending on the geometric relations among them, to form interpolation paths. This local mechanism, they argued, could represent the neural basis of IC interpolation. More general models built within similar, biologically-plausible frameworks have addressed a large range of related phenomena, including transparency and neon color spreading (Bressan, Mingolla, Spillman, & Watanabe, 1997; Grossberg & Mingolla, 1985). In the present work, we base our model on variations of the mechanisms and implementation developed in Heitger et al. (1992, 1998).

3. Motivation

Occluded contours are common in natural scenes, whereas ICs require carefully selected surface and background brightness and coloring (although they may be ecologically important in cases of camouflage). The striking similarity between the geometry that supports occluded and IC interpolation suggests that they share an underlying neural mechanism and that perhaps they have a common origin. The function of completing contours that appear partially occluded is clear – the contour fragments are perceptually completed because they are likely to actually be part of the same physical object. The identity hypothesis, and the model proposed here, is based in the idea that the visual system evolved to complete partially occluded contours and that the mechanism that does this also generates ICs.

Several physiologically consistent models of IC interpolation have been proposed, but none unifies illusory and occluded contour interpolation as posited by the identity hypothesis. A useful test of the identity hypothesis is to build a model based on the geometry of partial occlusion and then test it on both partial occlusion displays and displays for which subjects perceive ICs. The model presented here assesses the viability of local, bottom-up, image-based contour interpolation. The model’s performance when the input is that of an IC display will serve as a direct test of the idea that IC interpolation and occluded contour interpolation share a common initial process.

There were several motivating factors for using the work of Heitger et al. (1992, 1998) as the foundation for a unified model of both illusory and occluded contour interpolation. Two attractive features of the Heitger et al. (1998) model are that the underlying circuitry for calculation and representation is neurally plausible and the model in its full form is able to take grayscale images as input. The framework is thus capable of making both behavioral and physiological predictions.

What is more, while the calculations made by this model are inspired by the analysis of neurophysiological recordings, like the identity hypothesis it is specified exclusively as an algorithmic account of how the visual system interpolates contours. Following the three-level analysis framework of Marr (1982), we are making claims about interpolation primarily at the algorithmic or representational level. In terms of biological mechanisms, what we propose (as with Heitger et al., 1992, 1998) is broadly consistent with

known properties of cortical visual mechanisms, but is not meant to capture the full range of details as considered in recent physiological investigations or in some models of early visual processing (e.g., Grossberg, Mingolla, & Ross, 1997; Yazdanbakhsh & Grossberg, 2004).

Like its predecessors, the model we present here operates on local, non-symbolic representations of contour geometry and in turn generates local, non-symbolic descriptions of interpolated contours. Put less formally, the model generates a map of interpolation activation but does not ‘know’ that an interpolated edge and a continuous physically-specified edge all form part of the boundary of a single object. In object formation, early interpolation processes may provide inputs along with other scene constraints that ultimately lead to representations of objects, their arrangements, and their shapes (Kellman, Garrigan, & Shipley, 2005). Non-local, symbolic contour representations are very likely needed to explain perception of the unity and shapes of objects, and perhaps some aspects of the strength and specific shape of interpolated contours as well. A virtue of elaborating a model of the sort we present here is that such models may sharpen our understanding of what can and cannot be accomplished by relatively local, non-symbolic operations. The results can be used to infer what types of additional information must be incorporated into future models.

4. Grouping under occlusion

The computational core of our model is based on the framework defined in Heitger et al. (1998). In earlier work, Heitger et al. (1992) presented a biologically plausible approach to finding luminance-defined contours and contour junctions. We implemented and used the output of that model, the locations and orientations of edges and of L- and T-junctions, as the input to our interpolation mechanism. All inputs to the model, intermediate representations, and outputs from the model are specified in image coordinates that we call “maps”. Each map, therefore, is the size of the input image, and has values that code various information (e.g., the presence or absence of part of an interpolated contour) at each pixel location.

4.1. Model input

Our interpolation model requires as input two kinds of information: the location and orientation of contours and the location and orientation of contour endings (“key-points” formed at the intersection of contours that meet at the corners of individual objects and contours that meet when the edge of one object is occluded by another object). The Heitger et al. (1992) model provides a coarse-coded set of six orientation-selective contour maps (C maps, spaced at 60° increments) and twelve orientation-selective contour-end maps (ES maps, spaced at 30° increments). These maps provide all of the necessary information for our model to generate interpolated contour responses.

The ES maps generated from the Heitger et al. (1992) model are further processed into what will be denoted here as key-point maps (KP maps, Fig. 5C) by thresholding the spatially distributed responses of the ES maps. This reduces all contour-end responses that correspond to a single key-point into a discrete point at the same spatial location in all of the appropriate orientation channels. The relative magnitude of the responses in each of the twelve ES maps codes the precise orientation of the contours at the point where they meet.

Fig. 5B illustrates the summed C map responses across all orientations generated by processing the simple occlusion display shown in Fig. 5A. Individual points on each C map correspond to simulated complex cell responses at discrete locations and orientations. Note that the contour inputs to the model have a strength

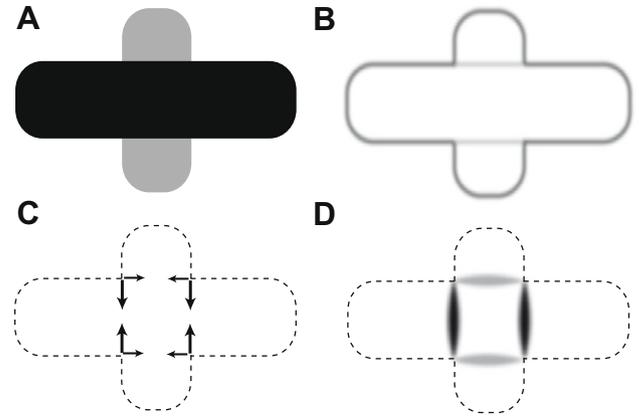


Fig. 5. An occlusion display, and the intermediary levels of processing within the model. (A) The initial occlusion display. (B) The C map response (designating contours), summed across all orientations. (C) KP map responses, illustrated as arrows corresponding to the magnitude and orientation of response. (D) τ grouping response.

(stronger = darker) that corresponds to the local contrast (Heitger et al., 1992). This dependency is maintained throughout, and affects the strength of computed interpolation paths.

The KP response locations and orientations are illustrated by arrows in Fig. 5C. Points on each KP map correspond to corners and contour junctions at discrete locations and orientations. An ideal corner of a square would therefore be represented as a point at the same spatial location on a KP map representing orientation θ , and a KP map representing orientation $\theta + 90^\circ$. There are half as many C maps as KP maps because the two contrast polarities (transitions of light \rightarrow dark vs. dark \rightarrow light) are combined in the C maps.

4.2. Detecting potential interpolation sites

In our model, the first step of the interpolation process is locating key-points. These are likely (but not guaranteed) to result from one surface occluding another. Next, τ , the T-junction map, marks the locations and orientations of T-junctions in visual scenes. Non-zero values of τ occur at locations on the image where a properly oriented key-point coincides spatially with a properly oriented contour:

$$\tau = KP \cdot \tilde{C} \quad (1)$$

The magnitude of τ for a specific orientation is the pixel-by-pixel product of a KP map with one orientation and an orthogonal \tilde{C} map. Since the maps are multiplied, activation must be present in both maps at the same location for generation of non-zero τ response. We use the notation \tilde{C} to refer to a range of oriented contour maps, in this case, $\pm 30^\circ$:

$$\tilde{C}^i = C^i + C^+ + C^- \quad (2)$$

where, for a specific orientation, the \tilde{C} map is the sum of the C map at one orientation (C^i) and the C maps at both neighboring orientations (C^+ and C^-). This allows the τ grouping function, G (defined below), to be sensitive to approximate T-junctions formed from the junction of contours that are not perpendicular to one another. For any specific orientation, the T-junction map is:

$$\tau^i = KP^- \cdot \tilde{C}^i \quad (3)$$

where KP^- corresponds to a KP map with orientation perpendicular to the orientation of \tilde{C}^i . Since the \tilde{C} maps are not sensitive to

contrast polarity, each \tilde{C} map will be paired with two different KP maps.

4.3. Interpolation responses

τ detects and represents the positions and orientations of T-junctions. Interpolation proceeds by comparing the locations and orientations of pairs of T-junctions to determine if they correspond to points of occlusion on a common contour. To achieve this, we use the same grouping-field algorithm defined in Heitger et al. (1998):

$$F(r, \theta) = \begin{cases} e^{-\frac{1-r^2}{\sigma^2}} \cdot \cos^{2n}(\theta) & \text{if } -\frac{\pi}{2} \leq \theta \leq \frac{\pi}{2} \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

We also use the same angular tuning parameter as in the earlier research ($n = 4$). This equation defines a set of locations, that we will call a “lobe” for the grouping mechanism. The example given is a lobe used to generate a grouping field oriented at 0° . At each point in the image, lobes are generated at 30° increments, corresponding to each of the 12 oriented τ maps.

Grouping occurs by first convolving combinations of specific lobes F^- and $\tau^{|\cdot|}$ maps thereby specifying all of the potential interpolation paths starting at every contour junction in the image. Interpolation requires contour junctions at both ends of the interpolation path. Relatable paths between the junctions are selected from the set of all possible interpolation paths:

$$G^- = \sqrt{(F^- * \tau^{|\cdot|}) \cdot (F^- * \tau^{|\cdot|})} \quad (5)$$

Response in the G map corresponds to the local representation of interpolation paths. Relatability is realized in the interpolation mechanism by only allowing specific combinations of τ and F . The multiplication of potential interpolation paths generates non-zero interpolation response in a grouping map G^- only when two τ responses are relatable. Example G map output collapsed across all orientations is illustrated in Fig. 5D. Since contour interpolation is not constrained to straight lines, curved interpolations are allowed by summing the activities of G response across multiple configurations. The total G response is defined as:

$$G = G_A + G_B + G_C + \frac{1}{2}(G_D + G_E + G_F + G_G) \quad (6)$$

where each of the subscripts corresponds to one of the configurations in Fig. 6. This determines the tolerance for curving interpolations, consistent with the geometric definition of contour relatability, which allows for curved interpolations up to about 90° .

The architecture described above is significantly simpler than the models on which it is based. The primary difference is that a single grouping operator in the current model replaces two distinct grouping operators (and a third operator that modulates between them) in the IC model described in Heitger et al. (1998). This change is motivated by the identity hypothesis and results in contour interpolation that is more consistent with human perception across a larger range of stimulus conditions.

Note that contour interpolation proceeds based on image features that are entirely specified by local contrast information. That is, features that require relational information, e.g., the angle at which two contours meet, are not available to the grouping operators. This leads to important ambiguities that have consequences for perceived interpolations. As we will demonstrate below, it is the lack of relational information that makes L- and T-junctions indistinct to local interpolation mechanisms and therefore allows the same mechanism to support both illusory and occluded contour interpolation.

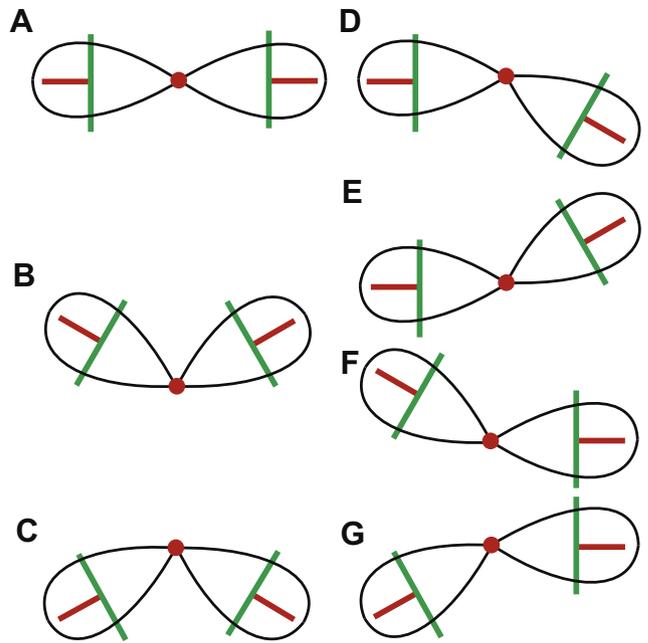


Fig. 6. The complete set of configurations that define G maps. For a given orientation (drawn above at 0°), these seven different lobe orientation combinations define the complete G map response. Configuration (A) interpolates collinear key-points, while configurations (B)–(G) interpolate with various types of curvature.

4.4. Global inhibition from ξ

After generating the set of G maps, we apply a global inhibitor ξ (Heitger et al., 1998):

$$\xi = 1 - \frac{\sum_{i=1}^{N/2} \sqrt{G_i \cdot G_{i\perp}}}{\sum_{i=1}^{N/2} (G_i + G_{i\perp})} \quad (7)$$

where N is the total number of G maps. ξ is also a map, bounded between zero and one. The value at each point on the map tends towards zero as the responses of spatially coincident orthogonal G maps increases. Each G map is multiplied by ξ , suppressing weaker activation when multiple crossing interpolation paths result from spurious grouping of coincidentally relatable T-junctions.

4.5. Final output

The model's final stage of processing takes both the original C maps and the ξ -thresholded G maps as input. For each of the channels, the C and G responses are added together into a common map. After superimposing image contours and interpolated contours, a spatial derivative is performed on each map at an orientation perpendicular to the orientation the map represents. This transformation results in an output map that contains discrete points of activation rather than the diffuse activation fields that the model has used up until this point. This output, presented in Section 5, below, is the collection of each of these output maps summed across orientations and displayed as an image.

5. Results

All simulations were executed on a Macintosh computer using MATLAB and the Image Processing Toolbox. In all output images, green points indicate C map activation, and red points indicate grouping map activation. Input images were 512×512 grayscale bitmaps.

5.1. Occluded contour interpolation

The model presented here is based in the framework of Heitger et al. (1998), but it has significantly different motivation and behavior. Heitger and colleagues proposed a model of IC interpolation. We have proposed a general model of contour interpolation consistent with the identity hypothesis. For this reason, we begin by illustrating how the current model performs the function for which it was designed, interpolation under conditions of partial occlusion, and how the model specified in Heitger et al. (1998) does not. Keep in mind that this demonstration does not indicate a failure of the Heitger et al. model because, as stated earlier, this model was not designed to interpolate occluded contours. Rather, these images are intended to illustrate the conditions under which interpolation initiated by τ is effective, and highlight the varied response of a different model built on a similar computational framework.

We implemented the model of Heitger et al. (1998) and compared its output to the output of our model. In Fig. 7, we compare contour interpolation initiated by τ response (G map activation) to the output generated by the Heitger et al. (1998) model for simple partial occlusion stimuli. The inputs to the model (Fig. 7, left column) are each perceived as three different configurations of two rectangles, a gray one above a black one. The human visual system interpolates behind the gray rectangle, joining the two, separated pieces of the black rectangle. It is evident that grouping of τ response (Fig. 7, middle column) yields robust contour interpolation for partial occlusion stimuli with contour junctions occurring at a

large range of angular relations. From top to bottom, occlusion of two rectangles oriented at 90° , 60° , and 30° relative to each other are shown. In the right column, the outputs of Heitger et al. (1998) is shown. From these three images, it is clear that the outputs of the Heitger et al. model are inconsistent for occlusion stimuli when contours meet at oblique angles.

5.2. Illusory contour interpolation

To test performance on stimuli for which ICs are perceived, the model performance was first evaluated using a standard Kanizsa triangle (Fig. 8A, left) as input. Observers viewing this stimulus typically report a white equilateral triangle in front of three black circles whose centers coincide with the vertices of the triangle. G and C map activation is shown in Fig. 8A, right. G map activation completes both the ICs (the sides of the triangle) and the occluded, circular inducing elements.

Another common figure used to demonstrate the perception of ICs is an Ehrenstein display (Ehrenstein, 1941). A standard, thin Ehrenstein display and a “thick” Ehrenstein display are shown in Fig. 8B, left and C, left, respectively. The input to the model for the thin Ehrenstein display was idealized as four line segments, each terminating at only one orientation exactly 90° relative to each of its neighbors. That is, the lines were represented with no thickness, and therefore, no corners. For this input, G map activation is weak and inconsistent with human perception (Fig. 8B, right).

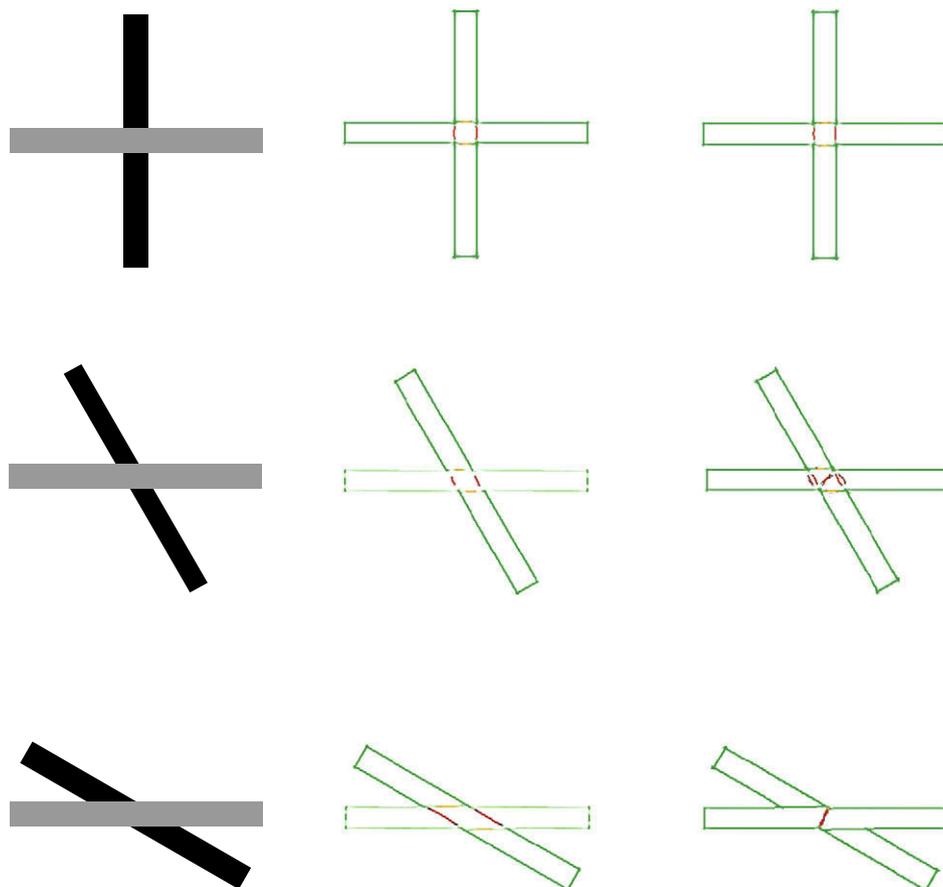


Fig. 7. Comparison of grouping initiated by τ response and the output of Heitger et al. (1998). Green points indicate C map activation and red points indicate grouping. Occlusion stimuli (left column) with varying angular configurations were used as input. The current model output (τ initiated grouping, middle column) is shown as well as the output of our replication of the Heitger et al. model (right column). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

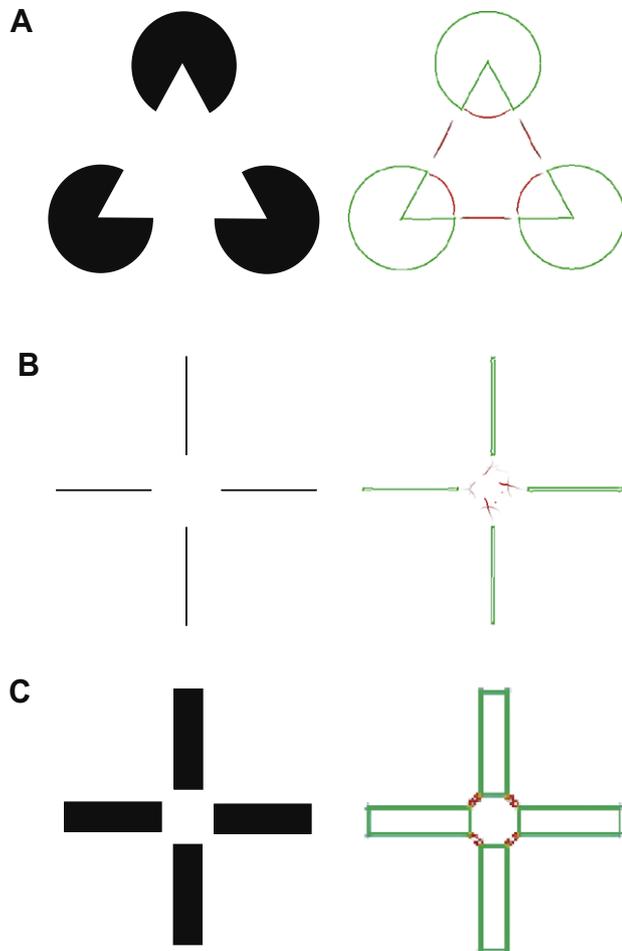


Fig. 8. Common IC illustrations. Input to the model (left) and *G* and *C* map activation (right) are shown. (A) Kanizsa triangle interpolation. *G* map activation completes both the triangle and the inducing elements demonstrating illusory and occluded contour interpolation respectively. (B) Thin Ehrenstein display interpolation. Importantly, the thin line segments that form the Ehrenstein display are represented as having unidirectional termination. That is, they are treated as line segments, not thin rectangles. *G* map activation is only low amplitude noise. (C) Thick Ehrenstein display interpolation. When the Ehrenstein display is composed of rectangles with bi-directional termination (corners), *G* map activation completes the central region, consistent with perceived IC interpolation.

In contrast, *G* map activation interpolates robustly for input from the thick Ehrenstein display (Fig. 8C, right). In this case, the input image was four rectangles oriented as in Fig. 8C, left. Each rectangle had two corners from which contour interpolation could initiate. The key difference between this figure and the thin Ehrenstein display is the presence of bi-directional contour terminations (corners) at the initiating sites of interpolation. We will argue below that this result as well as the weak, inconsistent *G* map activation resulting from the thin Ehrenstein display are consistent with an ecological interpretation of the geometry of thin lines.

5.3. The Petter effect

As discussed above, self-splitting objects, which give rise to crossing interpolations, have important implications for the identity hypothesis. Whereas completion for Petter effect displays is supported by our model (Fig. 9A), it is specifically prohibited by the model's forerunner (see Heitger et al., 1998, Fig. 5) because this is not an example of IC interpolation. In our model, L- and T-junctions are both suitable sites for initiation of contour interpolation. In our model, interpolation will occur, provided one of the contours

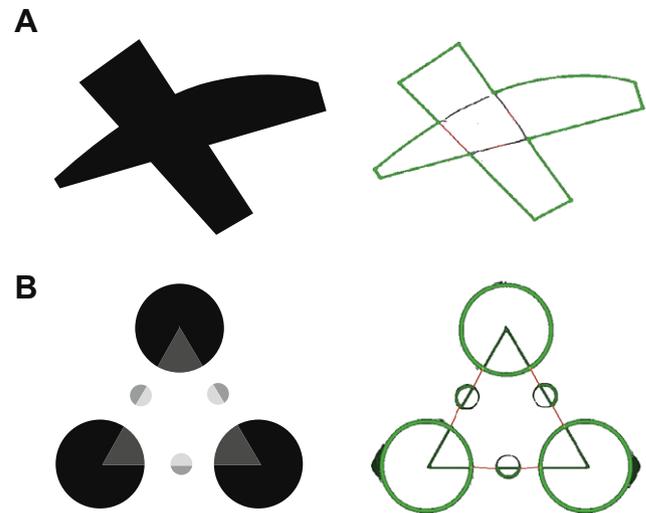


Fig. 9. Transparency and self-splitting objects. Input to the model (left) and *G* and *C* map activation (right) are shown. (A) Interpolation between all reliable L-junctions splits the figure into two parts, consistent with human perception. Assignment of modal and amodal status to the interpolations occurs at later stages of processing. (B) In transparency displays, like the one shown on the left, interpolation occurs between reliable T-junctions. The model correctly completes this image.

on each of the L-junctions form a reliable pair. Implications of this design choice are discussed below.

5.4. Transparency

Transparency displays, like the one shown in Fig. 9B, left, can be some of the most striking demonstrations of strong IC percepts. Note, however, that in the case shown, the ICs begin and end at L-junctions. Again, the current model's unified treatment of L- and T-junctions allow for contour interpolation to proceed in these cases (Fig. 9B, right).

5.5. Exotic cases

5.5.1. Quasimodal interpolation

Displays in which interpolation begins at a T-junction (and is amodal) and ends at an L-junction (and is modal) further stress the importance of treating L- and T-junctions equivalently at the earliest stages of contour interpolation. Fig. 10A, left is such a figure. It is apparent that the central region is perceived as a unified, white, quadrilateral, consistent with the model's interpolations (Fig. 10A, right).

5.5.2. Interpolation with mixed boundary assignment

G map activation was also tested for situations where ICs begin and end with inconsistent local boundary assignment. In Fig. 10B, left, the outside contours of the gray squares "belong" to the occluder whereas the inside boundaries of the notched circles "belong" to the occluded objects (the full circles partially occluded by the square). *G* map activation is insensitive to the local boundary assignment, and interpolates the contours of the square (Fig. 10B, right), consistent with human perception.

6. Discussion

The identity hypothesis is an ecologically appealing account of the primacy of occlusion in contour completion processes. In ordinary perception, amodal completion could operate in situations where one surface is in front of another, relative to the observer.

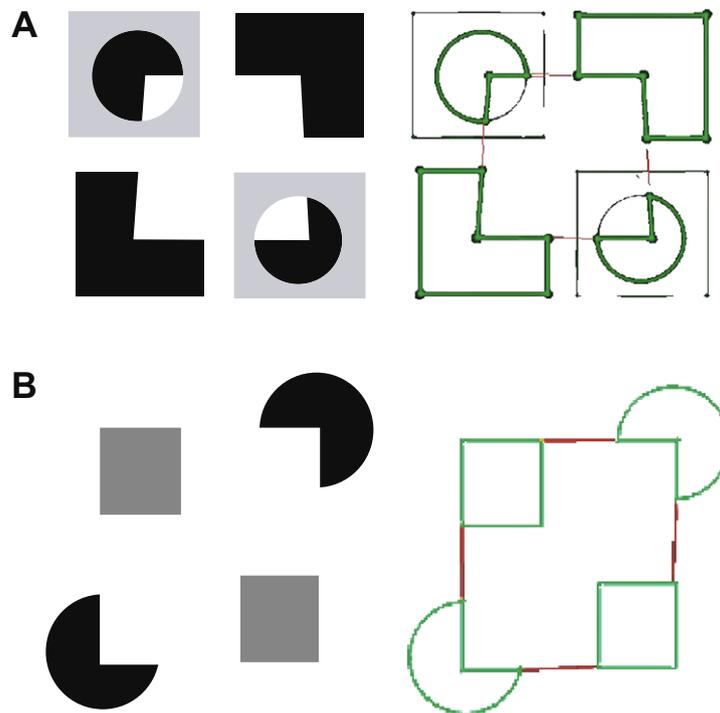


Fig. 10. Exotic cases of interpolation. Input to the model (left) and G and C map activation (right) are shown. (A) Quasimodal interpolation. In quasimodal interpolation, the interpolated part of each edge of the trapezoid begins at an L-junction (and is modal) and ends at a T-junction (and is amodal). (B) Interpolation with mixed boundary assignment. Interpolation can occur even when the boundaries of the to-be-completed object are defined as belonging to the interpolated object at one end of an interpolated contour, and belonging to an occluded inducing element at the other end. G map activation completes the central square under these conditions. The circle inducers are not completed, as the angle between key-points is 90° and beyond the limits of the tolerance for curved interpolation in the model.

Modal completion would operate for this same arrangement of surfaces, but only when the chromatic and luminance relations between the occluding surface and the background were carefully matched. When natural surfaces varying in lightness, color, and texture are considered, the latter case is a rarity. For this reason, the idea that a contour interpolation process designed for completing partially occluded contours also generates initial IC interpolations is, at the very least, parsimonious.

The model of contour interpolation proposed here provides a mechanism consistent with this account. The interpolation mechanism described above is designed to detect and interpolate between contour fragments using image features often indicative of partial occlusion (i.e. the presence of reliable T-junctions). The similarity of the initiating conditions for amodal and modal completion, especially at the local level at which this model operates, causes the model to also perform modal completion. The T-junctions that initiate occluded contour interpolation are similar enough to the L-junctions that initiate IC interpolation so that either junction type, in the proper arrangement, can act as input to a common interpolation mechanism. In our implementation, as in Heitger et al. (1998), we used only feed-forward, local processes. The model uses image features that tend to appear where contours become occluded, and it completes the portion of the contour that is not visible. We demonstrated the performance of the model on cases where observers report seeing the unity of partly occluded objects and on a number of figures for which subjects report perceiving ICs.

6.1. Simple illusory and occluded contour interpolation

It is evident from the output shown in Fig. 7 that τ initiated grouping successfully interpolates occluded contours across a large range of angular relationships at the T-junction. The results shown here indicate that simple occlusion events can be adequately re-

solved by local mechanisms simulating complex cells and a grouping process that incorporates contour reliability. The model also produces output consistent with IC perception. In Fig. 8A, grouping between τ response produces activation consistent with reported IC perception. Notably, τ grouping also completes the circular inducing elements.

Completion of both the illusory components (the sides of the triangle) and the occluded components (the notches in the circles) occurs despite the absence of T-junctions because the criteria for τ activation is satisfied by L-junctions as if they were weak T-junctions. T- and L-junctions are perceptually distinct, but the local oriented contrast associated with each is similar. τ activation marks the location of spatially coincident C-response and key-point response when the C-response is approximately perpendicular to the key-point response, as specified in Eq. (3). When this activation is present in each of a matched pair of the grouping fields (Eq. (4)), interpolation occurs.

Intuitively, τ response results from a contour ending at a location where another contour runs approximately perpendicular to it (a T-junction). At the level of local oriented contrast response this geometry is similar to two roughly perpendicular contours ending at the same location (an L-junction). L-junctions occur at the corners of objects and are also present at the points where IC interpolation initiates. To the model, L-junctions are encoded in a manner indistinguishable from weak T-junctions, and therefore support interpolation. They are weak because the C-response, integrated over the grouping field, is half as great. With L-junctions instead of T-junctions as the input, the same interpolation mechanism for occlusion events also produces ICs, consistent with the identity hypothesis.

The local, feed-forward design of this model demonstrates that interpolation can proceed (initially, at least), uninformed by higher-level, relational properties that may contradict the interpolation solution. While ignoring non-local and top-down constraints limits

what information is available at the interpolation stage (e.g. figure/ground assignment, border ownership), modeling interpolation in this way leads to a number of interesting predictions that mimic human perception.

One example is the Petter effect. Petter effect displays are perceived as two crossing surfaces despite the fact that there is no discontinuity of any surface properties across the entire surface (see Fig. 9A, left). From one perspective, the “simplest” interpretation of this figure is that there is only one object and that some of the contour discontinuities are formed at concave corners on the object. For human perceivers, however, there are two objects; one on top and completed modally, the other beneath and completed amodally. This is consistent with our model and with a single interpolation mechanism that accepts L-junctions as potential sites of modal or amodal completion. A stronger point can be made here, regarding the identity hypothesis. Although other types of models may be possible, within the style of the present model (and its predecessors), the identity hypothesis – an early interpolation process that underlies modal and amodal completion – must be embraced. The reason is that separate modal and amodal processes would need, for this kind of model, separate stimulus “triggers”, or different initiating conditions. This, indeed, was the intent of the model of Heitger and colleagues in using only L-junctions and attempting to account only for illusory, not occluded, contours. Self-splitting object displays, however, have identical junction information for both of the (crossing) interpolations; they are L-junctions in both cases. The appearance of one interpolation as modal and the other as amodal either shifts arbitrarily (for interpolations across relatively equal-sized gaps) or is determined by Petter’s Rule (for different-sized gaps). In either case, the modal/amodal appearance of each object is dependent on the other (one always appears in front, as modal). The implications of this interdependence for the identity hypothesis have been described previously (e.g., Kellman, Garrigan, & Shipley, 2005). The additional point made here is that filtering models of the sort used here cannot be built to have separate modal and amodal completion processes, because they would fail for Petter effect displays, where one modal and one amodal interpolation is produced from identical types of inducing features (specifically, L-junctions in a homogeneously colored region).

Another example of the evidence for unified modal and amodal contour interpolation comes from figures in which IC interpolation results in perceived transparency. Fig. 9B left, shows such a figure. Notice that here, modal completion is initiated by T-junctions, even though ICs are typically shown to begin at L-junctions (e.g., Fig. 8A, left), and T-junctions are typically associated with occlusion. Again, our model interpolates the contours in this figure appropriately (Fig. 9B, right), supporting a single contour interpolation mechanism for amodal and modal completion that is not determined by the presence of a particular junction type.

Another important consistency among the model, the identity hypothesis, and human perception is the example of quasimodal interpolation, where the interpolated boundary appears modal at some points along its extent and amodal elsewhere. In displays such as Fig. 10A, left, contour interpolation is initiated at a T-junction (and is amodal) and ends at an L-junction (and is modal). Again, the unified approach to contour interpolation taken by our model provides the correct interpolations (Fig. 10A, right). This display would not be predicted to produce significant interpolation responses in the model of Heitger et al. (1998), because the occluded inducer at one end would not activate the grouping operator in that model.

One last, informative example is illustrated in Fig. 10B. An important issue in final scene representations is border ownership: occlusion edges in scenes are boundaries of only one of the two surfaces that define the contour. Border ownership has sometimes

been argued to be an early input to interpolation processes, such that conflicts of border ownership might block interpolation from occurring in the first place (Anderson et al., 2002). It has been shown, however, that conflicting border ownership does not in general block contour interpolation (see for example, Fig. 16 in Kellman, Garrigan, & Shipley (2005)). Border ownership issues may weaken the appearance of edges in the final percept or may lead to percepts that are anomalous in terms of possible physical objects (Kellman, Garrigan, & Shipley, 2005), but agreement of border ownership does not appear to be a requirement for initial contour interpolation in the human visual system. The model proposed here fits with this observation, as shown in Fig. 10B. Here the contour begins as the edge of a square and ends as the occluding edge of a partly occluded circle. The physically specified stimulus is on one end figure and on the other end ground. The existence of mixed-boundary interpolation is a specific prediction of the identity hypothesis, consistent with perceived interpolated contours, and consistent with our model’s behavior.

6.2. Limitations

So far we have shown that a local, feed-forward model of contour interpolation properly completes contours under a range of basic illusory and occluded contour interpolation conditions. In this section we discuss a number of conditions in which the model output is not consistent with human perception.

Our model does not generate interpolated contours consistent with human perception for the classic Ehrenstein display, as in Fig. 8B. Specifically, this is the case if the thin line-ends of the Ehrenstein display are represented as ideal one-dimensional mathematical objects, that is, contours extending in one direction only, without corners. When presented with this input, τ -initiated grouping first joins the occluded line inducers as thin crossing bars. This grouping is then suppressed by a thresholding function which attenuates crossing interpolated contours (referred to and defined in Heitger et al., 1998 as ζ). The function of the ζ map is to suppress weak interpolation responses that likely arise from distal key-points with coincidentally relatable orientation. The resulting output after thresholding is simply low amplitude noise, drawn in red in Fig. 8B, right.³

The “lines” of the Ehrenstein display are not ideal, one-dimensional objects, of course; physically they are thin rectangles. The lack of corner detection by our model is an issue of resolution. As we move away from the ideal (but degenerate) case of very thin lines without perceivable corners, the model performance changes significantly. Once the inducers are increased in width so that corners of each inducer are represented, interpolation between τ responses correctly completes the contours (Fig. 8C). This limitation suggests that corners are represented even when the physically specified stimulus elements are at a scale too small for the corners to be perceptually salient.

Another limitation of the local computations underlying this model is that the orientations of the contours that meet to form L-junctions are ambiguous at the point where the two contours meet. This can cause the model to generate identical key-point response from very different contour configurations. In Fig. 11A, two properly oriented acute triangles can fool the model into responding as if the input were a single, horizontal line occluded by a rectangle oriented perpendicular to it. This happens because the contours defining the triangles are oriented so that the key-point response from each edge is represented evenly across two adjacent orientation maps. Because key-point responses at single locations

³ For interpretation of color in Figs. 1–6, the reader is referred to the web version of this article.

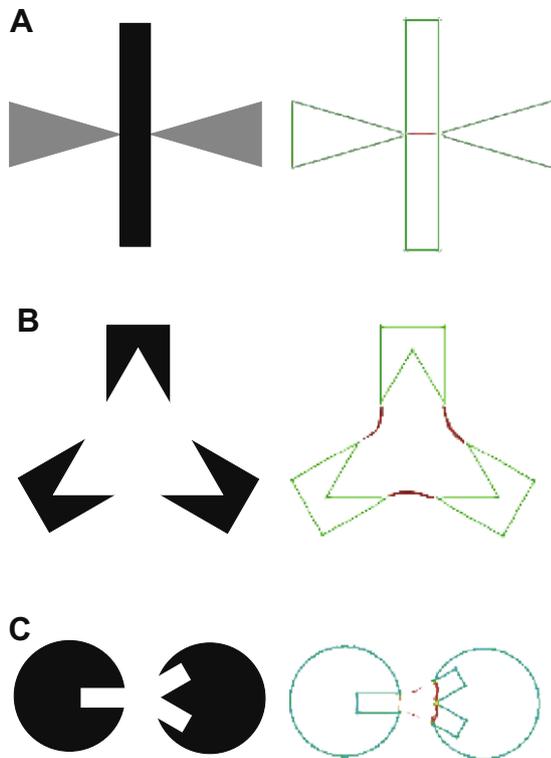


Fig. 11. Limitations of local operations. Input to the model (left) and *G* and *C* map activation (right) are shown. (A) Spurious interpolation. Superposition of key-point responses fools the model into generating output as if the tips of each triangle were lines terminating perpendicular to the occluder. (B) Effects of incidental geometry on model output. The angle where the inducing element edge intersects the illusory figure affects the shape of the interpolation paths. (C) An illustration of promiscuous grouping. Because of the local nature of the computations underlying grouping, key-points that correspond to a single contour in the world can fork and become anchors for multiple simultaneous ICs. This suggests that a complete account of interpolation must ultimately be constrained to join two key-points uniquely, consistent with a physically possible description of the world.

are averaged, the two edges of each triangle create a key-point response with highest magnitude at the orientation directly between each of the adjacent orientation maps. The response is very similar to the response that results from an occluded line. Though conditions that generate this form of erroneous output will always occur in a model of this architecture, more finely-tuned coding of orientation (more than six *C* maps and twelve *KP* maps) could minimize this effect. With a sufficient number of orientation maps, the model would be able to correctly distinguish between acute angles in a manner consistent with human perception.

The ambiguous orientation at the point where contours meet presents another problem, however, that cannot be resolved with a larger set of orientation channels. Averaging local key-point response causes the paths of interpolation to be influenced by the orientations of both contours at a junction, even though only one contour's orientation is relevant. The shapes of perceived ICs are not affected by the relative orientations of incidental contours ending at an L-junction (Fig. 11B, left), but the model output (Fig. 11B, right) is altered. The model cannot associate part of the activation at a key-point with the corresponding contour that generated it. At an L-junction, two key-points are generated at the same spatial location for each of the contours ending there. Because there is no mechanism in this model for binding individual contours to particular components of *KP* response, there is no way to distinguish the origins of grouping responses in the *G* maps. The grouping activation from independent contours is conflated. When the *G* maps are reduced to the final interpolation path, as shown in Fig. 11B,

the path is a distorted average of the two interpolations the figure supports (one interpolation completing the triangle, and another connecting the relatable inducing rectangles). To overcome this limitation, it would be desirable to associate each contour that leads into a junction with a specific interpolation path. In principal, *KP* maps could be associated to individual contours, and this segregation could persist through the generation of *G* maps and beyond. This is a significant departure from the current architecture, however, and indicates that contours may need to be represented symbolically prior to interpolation.

There are also instances in which processes not possible within the architecture of this model may be necessary to assess the outcomes of contour interpolation or allow for competition among several possible “perceived” figures. In Fig. 11C, left, a single rectangle partially occludes the black circle on the left, and two rectangles partially occlude the black circle on the right. The percept is usually a Y-shaped occluding object, although two individual completions of the left occluded rectangle, each as a longer, bent rectangle are also possible. Our model cannot choose among the multiple interpolation solutions. (Of course, phenomenology also appears to have some difficulty in choosing, as all of these appearances are possible and appear at various times with prolonged viewing.) The model output, in Fig. 11C, right, has weak interpolation along the horizontal connections and misshapen completion of the right circle. Once again, this difficulty has been produced by making the edges leading into junctions converge at small angles, revealing the limitations of averaging and suppression in the model.

The simplest additional constraint that could be added to the model would be encoding from ensembles of activated oriented units the unique orientation of each edge leading into a tangent discontinuity. This information is already implicitly contained in the model (and in Heitger et al., 1992, see Fig. 10). We have stopped short of this step because it involves an explicit transition from use of non-symbolic activations of detectors to symbolic coding of orientation. To reveal both the capabilities and limitations of non-symbolic models, we have confined the current work to the former, leaving addition of higher level constraints for future work.

Among other higher level constraints that may combine with outputs of the present model to determine final scene perception are consistency of boundary assignment and closure (whether real and interpolated boundaries together define a closed region). As presented, the model generates ICs originating at L-junctions if their local activation patterns in the key-point and *C* maps are sufficiently similar to those produced by T-junctions. The model interpolates between all properly configured L-junctions, even though many L-junctions that are relatable do not support clear IC formation. In Fig. 12A, top, L-junctions generate ICs. However, in Fig. 12A, bottom, a similar pair of L-junctions give rise to, at best, relatively weak ICs. Both cases would, if presented to our model, generate the same relatively strong ICs. Non-local processes, such as consistency of border ownership or enclosure of a surface region may be necessary for explaining these phenomenological differences.

This division of labor – early, promiscuous interpolation plus the operation of other scene constraints – is supported by behavioral studies of normal human perception (Kellman et al., 2007) and by comparisons of healthy human observers and subjects with a particular visual deficit called “integrative agnosia” (Riddoch & Humphreys, 1987). Guttman and Kellman (2001) used displays with relatable edges but weak apparent contours (e.g., IC displays made with outline figures) to show evidence for early interpolations between the relatable edges, despite the absence of clear phenomenal contours in the final scene representation. Evidence from patients with integrative agnosia, a disorder in which patients have difficulty binding elements into configural wholes, supports this conclusion. Specifically, Giersch, Humphreys, Boucart, and Kovacs

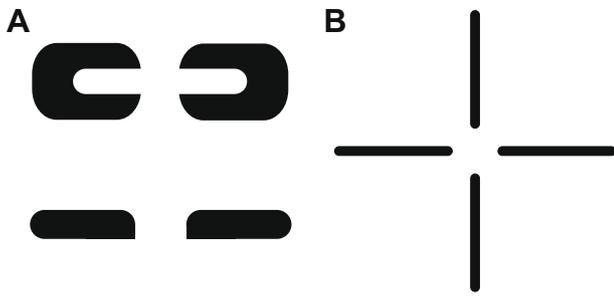


Fig. 12. Global considerations. (A) The need for global constraints on interpolation. For input like the top shape, our model would correctly interpolate two ICs. Because the model is restricted to local computations, however, spurious grouping can also occur whenever the geometry of a pair of key-points is relatable, even if global geometric constraints suppress the illusory percept. For input like the bottom shape, our model would generate an IC equal in strength to the corresponding IC in the top display, despite the relatively weak, or absent, IC percept generated by such a figure. (B) An Ehrenstein display with rounded corners is shown. Without the presence of key-points, the illusory percept is completely absent. This indicates that the visual system is processing the thin lines in the canonical Ehrenstein display as having corners, even if they are perceptually ambiguous due to resolution limitations.

(2000) found that one patient, HJA, had impairment in visual judgments involving occluded figures when the occluded figures were more likely to be perceptually completed (e.g., the interpolation path was short). HJA appeared to have intact interpolation mechanisms, but impaired ability to form coherent shapes from the representations of visible and interpolated parts. Behrmann and Kimchi (2003) found similar results in two other patients also presenting symptoms of integrative agnosia. Giersch, Humphreys, Barthaud, and Landmann (2006) have further shown that similar configural deficits can be elicited in healthy patients given Lorazepam, a benzodiazepine drug that interacts with GABA at the GABA_A receptor, and is also known to impair the integration of visual features into configurations (Giersch, 1999).

6.3. Contrast polarity of input edges

A recent report suggests that the angular range of contour interpolation depends on the contrast polarities of edges leading into occlusion (Geisler & Perry, 2009). Consistent with evidence that interpolation occurs between inputs of opposite contrast polarity (Field et al., 1993; Kellman & Loukides, 1987; Kellman et al., 2007; Spehar & Clifford, 2000; Victor & Conte, 2000), our model is largely insensitive to the contrast polarity of relatable edges. Based on scene statistics, Geisler and Perry (2009) argued that opposite polarity contour fragments are unlikely to be connected unless they are nearly collinear (within about 5–10°). They also reported that subjects' subjective impressions about interpolation correlated with scene statistics.

Why the scene statistics suggesting that edge fragments of opposite contrast polarity are less likely to be part of the same contour is unknown. Perhaps scenes in which edge fragments satisfy certain geometric relations but have opposite contrast are more likely to involve more complex occlusions, junction structures, or presence of more edges of different surfaces within a local region than those with uniform contrast polarity.

Implications of these observations about scene statistics for interpolation models is not clear. When using the prior measured from their scene statistics, Geisler and Perry (2009) note that the ideal rule for interpolation connects virtually no elements of opposite contrast polarity and even for edge elements of the same polarity, it includes only elements that are within a few degrees of collinear (see Fig. 7b in Geisler & Perry (2009)). This result is contradicted by studies showing that interpolation has considerable

tolerance for curvature (Field et al., 1993; Field, Hayes, & Hess, 2000; Fulvio, Singh, & Maloney, 2008; Shipley & Kellman, 1992a). Geisler and Perry (2009) obtained a more plausible ideal observer by removing the dependence of interpolation on distance and setting the prior that any two edge fragments are connected to 0.5. This approach was also applied to their experiment: subjects were instructed that half of the pairs connected behind the occluder and half did not. Obtaining subjective judgments under such a constraint is not the same as assessing ordinary perception of connectedness. In particular, instructing subjects that half of all pairs connect may introduce demand characteristics (especially in a study in which half of edge pairs had the opposite and half had the same contrast polarity).

Without introducing a new prior not derived from the scene statistics, these results are not in accord with human interpolation performance. Specifically regarding contrast polarity, Field et al. (2000) found reduced overall performance in their path detection task with opposite polarity (reversed phase Gabors), but the effect of angular relations appeared to be independent of contrast polarity. We interpret these results as indicating that a similar geometry underlies contour linking performance for both same and opposite contrast polarities, with some damping of overall performance in the latter case. Such a pattern may be consistent with the operation of both phase-sensitive and phase-insensitive contour integration mechanisms having somewhat different processes (Field et al., 2000) and/or with early contour linking mechanisms whose contribution to contour representations and phenomenology depend on additional scene constraints (Kellman, Garrigan, & Shipley, 2005).

One important way in which the model presented here and related models (Kellman & Shipley, 1991; Kellman et al., 2007) differ from the approach of Geisler and Perry (2009) is that these models have constraints besides relatability, specifically the importance of junction structures ("key-points" or tangent discontinuities) in triggering interpolation. The restricted geometry of interpolation and its interaction with contrast in the Geisler and Perry analysis may reflect information that is not included in their model, namely the arrangements of tangent discontinuities (key-points), intervening (crossing) edges (Kellman et al., 2007), and support ratio (Banton & Levi, 1992; Shipley & Kellman, 1990). Were these features taken into account, presumably more candidate edge elements would belong to the same edge contour, and the derived decision criterion would more closely match performance.

Deriving a model of contour interpolation that includes these effects from scene statistics is difficult because the corresponding distributions are hard to sample. The approach presented here for extracting and categorizing the junctions in an image could, in theory, be combined with scene statistics and used to generate an idealized grouping mechanism. While this would require significant additions to the current model, the use of image statistics could provide a valuable, quantitative comparison to the current model's performance on standard occlusion displays, and would help reveal the connection between the constraints embodied in the current model and the usefulness of interpolation processes in coping with the statistics of the environment.

6.4. Neural locus of contour interpolation

Whereas the model described here relies entirely on local, retinotopic image processing, there are at least two reasons to believe that higher-level regions of visual cortex must also be involved in contour interpolation prior to the formation of representations of surfaces and objects. One reason is that quantities that cannot be measured locally on an image (e.g., support ratio; Shipley & Kellman, 1992b) affect interpolation processes. Another reason is that contour interpolation has been shown to operate in 3D (Kellman, Garrigan, Yin, et al., 2005), even when depth information is

presented only through stereo disparity signals that are not believed to be available in early visual encodings (Cumming & Parker, 1999). Higher-level regions that appear to encode orientations and position in 3D space independent of how these orientations and positions are specified (e.g., cIPS; Sakata, Taira, Kusunoki, Murata, & Tanaka, 1997) could support 3D interpolation.

The importance of higher-level visual areas does not, however, mean that lower-level regions are unimportant for contour interpolation. Evidence for unified representation of illusory and occluded contours have been found in low-level visual areas (e.g., cell recordings in monkey V1/V2; Bakin et al., 2000; Sugita, 1999) and in higher-level encodings (e.g., ERP mappings in LOC; Murray et al., 2004). Murray et al. (2004) suggested that an interaction of low and high-level visual processing may be at work in the formation of interpolated contours, and that equivalent activation in low level regions to both illusory and occluded contours may reflect feedback from higher-level regions. Indeed, the general consensus from the neuroimaging literature measuring human observers suggests that the formation and representation of illusory contours recruits a diverse network of cortical foci, though the specific circuits are still a matter of scientific inquiry (for a detailed review, see Seghier & Vuilleumier (2006)). This raises the question of where in this more complicated neural picture does our model fit in? Our model appears to be consistent with a low-level, image-driven component of a contour interpolation mechanism that spans several levels of visual processing. The complementary, high-level symbolic component of this mechanism is not specified here, and connecting these early and later visual encodings remains one of the fundamental challenges of visual neuroscience.

6.5. Relation to other models

The computational architecture on which our model is built is based on the framework described in Heitger et al. (1998). The work detailed in the current paper is consistent with other models of perceptual processes, but adds an important contribution. Perhaps most important, our model provides a contour interpolation mechanism consistent with the identity hypothesis, illustrating the general feasibility and providing a specific instantiation of an early contour interpolation stage underlying modal and amodal completion. Such a mechanism could be realized within other models, such as the Boundary Contour System of the FACADE theory (Grossberg & Mingolla, 1985) a general purpose model that has been applied to many perceptual phenomena.

Our approach is complementary to modeling contour interpolation from a representational perspective. For example, Ullman (1976) first proposed that contour interpolations could be modeled as a path between two contour endpoints consisting of two smoothly-joined, circular arcs tangential to the endpoints with the lowest total curvature. A different representation, involving a line segment and a circular arc is derived in Kellman and Shipley (1991). Other computational accounts of interpolated contour representations have been suggested (e.g., Williams & Jacobs, 1997), but currently there is little behavioral evidence to determine which best describes perception. This problem is complicated by the relative similarities of the contour shapes described by each method (relative to the variability of typical behavioral data).

7. Conclusion

The model presented here robustly completes contours through a variety of contour geometries under conditions of partial occlusion. When illusory figures are provided as input to this occlusion model, we were also able to replicate a variety of perceptual phe-

nomena, from standard IC completion (Fig. 8A) to less intuitive contour completion with mixed boundary assignment (Fig. 10B). Our model adds substantive improvements to its predecessor (Heitger et al., 1998). In the largest perspective, it is more ecologically motivated than any model attempting to develop a mechanism specifically for interpolating ICs separately from occluded contours. Interpolation mechanisms seem more likely to have evolved to handle partial occlusion, which exists pervasively in natural scenes, whereas ICs are found in abundance primarily within psychology textbooks.

The computations by which contour interpolation occurs in our model are also less complicated than the corresponding calculations of the Heitger et al. (1998) model. In order for their model to account for different classes of illusory stimuli, two separate mechanisms were posited. One interpolates between thin line ends, like those defining the standard Ehrenstein display. The other interpolates between corners like those in a Kanizsa triangle. These two mechanisms compete through an additional operator that is sensitive to the configuration of the key-points at the sites of interpolation. Our proposed model needs no such modulation: all interpolation proceeds through a single unified operator.

As reported earlier in this paper, our model has several limitations that give insight into higher-level, non-local properties that also influence perception of interpolated contours. The first limitation is that our model is unable to correctly interpolate a traditional Ehrenstein display (Fig. 8C). This deficiency only occurs in the pathological case where the inducing elements of the display are represented as lines, with zero width. This raises the question as to whether or not it is reasonable to expect the visual system to represent an object in such a physically-impossible way, simply because a feature that can be inferred (the presence of a corner) is below the resolution of the system. Clearly if the object exists in the world, it must have some non-zero width. By zooming in on the display, effectively making the inducers more obviously rectangular, our model is able to correctly interpolate and complete the illusory figure, while the Heitger et al. (1998) model does not generate response consistent with human perception.

If we consider the thicker Ehrenstein display (Fig. 8C), ICs are perceived, and interpolation is generated by our model. This approach is also compatible with psychophysical evidence indicating that enlarging, such inducers but giving them perceptibly rounded tangent discontinuities, weakens or eliminates interpolation (Shipley & Kellman, 1990). Fig. 12B shows one example. Without corners, no key-points are generated, and our model will also not generate interpolated contours with this stimulus as input. This is further evidence for the idea that as the width of the inducers decreases to the point where the presence of corners becomes ambiguous, the visual system assumes that corners are present.

Our model cannot make figure-ground assignments, unlike some other models (Grossberg & Mingolla, 1985; Heitger et al., 1998). Contour interpolation in cases of ambiguous boundary assignment, as in Fig. 10B, however, suggest that some figure-ground assignments remain undefined at this stage of processing. This does not mean that all aspects of border ownership are decided after interpolation; that is likely not the case, especially when the best sources of depth order information (such as binocular disparity or accretion-deletion of texture) are present. Still, as shown in Fig. 10B and elsewhere, early commitment regarding border ownership is inconsistent with human interpolation phenomena.

Also, contours created by interpolation itself may pose new issues of border ownership and may create units for which figure-ground assignment must be made. Perhaps the most familiar example is that interpolation in a standard Kanizsa triangle changes figure-ground assignment for some boundaries as a consequence of interpolation. (A partial-circle inducing element

in isolation would “own” all of its contours, but the Kanizsa triangle owns some of them when the element is part of an interpolation display.) Some calculations establishing figure–ground assignment and consistency of border ownership may be part of later surface processing and integration stages. These stages also may have the final say on whether interpolated contours have modal or amodal appearance in the ultimate scene representations (Kellman et al., 1998, 2005).

The most informative aspects of our model may be its limitations. These limitations are not unique to our model, but instead challenge the class of contour interpolation models that relies on purely local computations. Specifically, without some explicit knowledge of the contours that gave rise to the pattern of activation at a key-point, the shape of interpolation can be biased by incidental geometry (as in Fig. 11A), or entirely inappropriate (as in Fig. 11B). To overcome these limitations, contours leading into key points must be coded in a manner that establishes their orientation and keeps them independent of other contours in the scene. The interpolation mechanism can then use this information to establish illusory and occluded contours that are consistent with the relevant geometry and unaffected by nearby incidental features. By coding key-point response in such a manner, information about the length of the terminating contour also becomes available. This contour length information could then be used as the quantity that determines the magnitude of the key-point response, which would be more consistent with quantitative variation in interpolation strength (Shipley & Kellman, 1992b).

Overall, the success of our model in unifying illusory and occluded contour interpolation suggests that contour interpolation of both types likely arises from a single mechanism at the earliest stages of processing. Success with this broader spectrum of interpolation phenomena suggests that initial interpolation across gaps based primarily on basic filtering operations, as originally suggested by Heitger et al. (1998), by Grossberg and colleagues, and by others, is likely a fundamental, early step along the way to object formation. The limitations of our model indicate that at later stages, non-local processing of image structure is necessary for fully explaining formation of perceived objects and their relations to other surfaces in a scene.

References

- Albert, M. (2001). Surface perception and the generic view principle. *Trends in Cognitive Sciences*, 5(5), 197–203.
- Albert, M. (2007). Mechanisms of modal and amodal interpolation. *Psychological Review*, 114(2), 455–469.
- Albert, M., & Hoffman, D. (2000). The generic-viewpoint assumption and illusory contours. *Perception*, 29(3), 303–312.
- Anderson, B. L. (2007). The demise of the identity hypothesis and the insufficiency and nonnecessity of contour relatability in predicting object interpolation: Comment on Kellman, Garrigan, and Shipley (2005). *Psychological Review*, 114(2), 470–487.
- Anderson, B. L., Singh, M., & Fleming, R. W. (2002). The interpolation of object and surface structure. *Cognitive Psychology*, 44, 148–190.
- Bakin, J. S., Nakayama, K., & Gilbert, C. (2000). Visual responses in monkey areas V1 and V2 to three-dimensional surface configurations. *Journal of Neuroscience*, 21, 8188–8198.
- Banton, T., & Levi, D. M. (1992). The perceived strength of illusory contours. *Perception & Psychophysics*, 52(6), 676–684.
- Baylis, G. C., & Driver, J. (1993). Visual attention and objects: Evidence for hierarchical coding of location. *Journal of Experimental Psychology: Human Perception & Performance*, 19(3), 451–470.
- Behrmann, M., & Kimchi, R. (2003). What does visual agnosia tell us about perceptual organization and its relationship to object perception? *Journal of Experimental Psychology: Human Perception & Performance*, 29(1), 19–42.
- Bressan, P., Mingolla, E., Spillman, L., & Watanabe, T. (1997). Neon color spreading: A review. *Perception*, 26(11), 1353–1366.
- Buffart, H., Leeuwenberg, E., & Restle, F. (1981). Coding theory of visual pattern completion. *Journal of Experimental Psychology: Human Perception & Performance*, 7, 241–274.
- Corballis, P. M., Fendrich, R., Shapley, R., & Gazzaniga, M. S. (1999). Illusory contours and amodal completion: Evidence for a functional dissociation in callosotomy patients. *Cognitive Neuroscience*, 11, 459–466.
- Cumming, B. G., & Parker, A. J. (1999). Binocular neurons in V1 of awake monkeys are selective for absolute, not relative, disparity. *Journal of Neuroscience*, 19, 5602–5618.
- Ehrenstein, W. (1941). Über abwandlungen der L. Hermannshen helligkeitserscheinung. *Zeitschrift für Psychologie*, 150, 83–91.
- Fantoni, C., & Gerbino, W. (2003). Contour interpolation by vector-field combination. *Journal of Vision*, 3(4), 281–303.
- Fantoni, C., & Gerbino, W. (2008). Interposition, minimal depth, and depth-from-disparity. *Journal of Vision*, 3(4), 853a.
- Field, D., Hayes, A., & Hess, R. F. (1993). Contour integration by the human visual system: Evidence for a local “association field”. *Vision Research*, 33, 173–193.
- Field, D., Hayes, A., & Hess, R. F. (2000). The roles of polarity and symmetry in the perceptual grouping of contour fragments. *Spatial Vision*, 13(1), 51–66.
- Fulvio, J., Singh, M., & Maloney, L. (2008). Precision and consistency of contour interpolation. *Vision Research*, 48(6), 831–849.
- Geisler, W. S., & Perry, J. (2009). Contour statistics in natural images: Grouping across occlusions. *Visual Neuroscience*, 26, 109–121.
- Geisler, W. S., Perry, J., Super, B., & Gallogly, D. (2001). Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41, 711–724.
- Giersch, G. (1999). A new pharmacological tool to investigate integration processes. *Visual Cognition*, 6, 267–297.
- Giersch, G., Humphreys, G., Barthaud, J., & Landmann, C. (2006). A two-stage account of computing and binding occluded and visible contours: Evidence from visual agnosia and effects of lorazepam. *Cognitive Neuropsychology*, 23(3), 261–277.
- Giersch, G., Humphreys, G., Boucart, M., & Kovacs, I. (2000). The computation of occluded contours in visual agnosia: Evidence for early computation prior to shape binding and figure–ground coding. *Cognitive Neuropsychology*, 17(8), 731–759.
- Gold, J. M., Murray, R. F., Bennett, P., & Sekuler, A. (2000). Deriving behavioral receptive fields for visually completed contours. *Current Biology*, 10, 663–666.
- Grossberg, S. (1994). 3-D vision and figure–ground separation by visual cortex. *Perception & Psychophysics*, 55, 48–120.
- Grossberg, S., & Mingolla, E. (1985). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review*, 92, 173–211.
- Grossberg, S., Mingolla, E., & Ross, W. D. (1997). Visual brain and visual perception: how does the cortex do perceptual grouping? *Trends in Neuroscience*, 20(3), 106–111.
- Guttman, S., & Kellman, P. J. (2001). Contour interpolation: Necessary but not sufficient for the perception of interpolated contours. *Journal of Vision*, 1(3), 384a.
- Guttman, S., & Kellman, P. J. (2004). Contour interpolation revealed by a dot localization paradigm. *Vision Research*, 44, 1799–1815.
- Heitger, F., Rosenthaler, L., von der Heydt, R., Peterhans, E., & Kubler, O. (1992). Simulation of neural contour mechanisms: From simple to end-stopped cells. *Vision Research*, 23, 963–981.
- Heitger, F., von der Heydt, R., Peterhans, E., Rosenthaler, L., & Kubler, O. (1998). Simulation of neural contour mechanisms: Representing anomalous contours. *Image and Vision Computing*, 16, 407–421.
- Hirsch, J., De La Paz, R. L., Relkin, M. R., Victor, J., Kim, K., Li, T., et al. (1995). Illusory contours activate specific regions in human visual cortex: Evidence from functional magnetic resonance imaging. *Proceedings of the National Academy of Science*, 92, 6469–6473.
- Kellman, P. J. (2003). Segmentation and grouping in object perception: A 4-dimensional approach. In M. Behrmann & R. Kimchi (Eds.), *Perceptual organization in vision: Behavioral and neural perspectives: The 31st carnegie symposium on cognition*. Hillsdale, NJ: Erlbaum.
- Kellman, P. J., Garrigan, P., Kalar, D., & Shipley, T. (2003). Good continuation and relatability: Related but distinct principles. *Journal of Vision*, 3(9), 120a.
- Kellman, P. J., Garrigan, P., & Shipley, T. F. (2005). Object interpolation in three dimensions. *Psychological Review*, 112(3), 586–609.
- Kellman, P. J., Garrigan, P., Shipley, T. F., & Keane, B. (2007). Interpolation processes in object perception: A reply to Anderson. *Psychological Review*, 114(2), 488–502.
- Kellman, P. J., Garrigan, P., Yin, C., Shipley, T. F., & Machado, L. (2005). Three-dimensional interpolation in visual object recognition: Evidence from an objective performance paradigm. *Journal of Experimental Psychology: Human Perception & Performance*, 31(3), 558–583.
- Kellman, P. J., Guttman, S. E., & Wickens, T. D. (2001). Geometric and neural models of object perception. In T. F. Shipley & P. J. Kellman (Eds.), *From fragments to objects: Segmentation and grouping in vision*. Oxford, UK: Elsevier Science.
- Kellman, P. J., & Loukides, M. G. (1987). An object perception approach to static and kinetic subjective contours. In S. Petri & G. E. Meyer (Eds.), *The perception of illusory contours*. New York: Springer.
- Kellman, P. J., & Shipley, T. F. (1991). A theory of visual interpolation in object perception. *Cognitive Psychology*, 23, 141–221.
- Kellman, P. J., Yin, C., & Shipley, T. F. (1998). A common mechanism for illusory and occluded object completion. *Journal of Experimental Psychology*, 24(3), 859–869.
- Lerner, Y., Hendler, T., & Malach, R. (2002). Object-completion effects in the human lateral occipital complex. *Cerebral Cortex*, 12, 163–177.
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Mendola, J. D., Dale, A. M., Fishl, B., Liu, A. K., & Tootell, R. B. (1999). The representation of illusory and real contours in human cortical visual areas revealed by functional magnetic resonance imaging. *Journal of Neuroscience*, 19, 8560–8572.

- Michotte, A., Thines, G., & Crabbe, G. (1964). *Les complements amodaux des structures perceptives [Amodal complements and perceptual organizations]*. Louvain, Belgium: Universitaires de Louvain.
- Murray, M. M., Foxe, D. M., Javitt, D. C., & Foxe, J. J. (2004). Setting boundaries: Brain dynamics of modal and amodal illusory shape completion in humans. *Journal of Neuroscience*, 24(31), 6898–6903.
- Petter, G. (1956). Nuove ricerche sperimentali sulla totalizzazione percettiva. *Revista di Psicologia*, 50, 213–227.
- Riddoch, M., & Humphreys, G. (1987). A case of integrative visual agnosia. *Brain*, 110, 1431–1462.
- Ringach, D. L., & Shapley, R. (1996). Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision Research*, 36, 3037–3050.
- Sakata, H., Taira, M., Kusunoki, M., Murata, A., & Tanaka, Y. (1997). The TINS lecture. The parietal association cortex in depth perception and visual control of hand action. *Trends in Neurosciences*, 20(8), 350–357.
- Seghier, M., & Vuilleumier, P. (2006). Functional neuroimaging findings on the human perception of illusory contours. *Neuroscience and Biobehavioral Reviews*, 30, 595–612.
- Shipley, T. F., & Kellman, P. J. (1990). The role of discontinuities in the perception of subjective figures. *Perception & Psychophysics*, 48, 259–270.
- Shipley, T. F., & Kellman, P. J. (1992a). Perception of partly occluded objects and illusory figures. *Journal of Experimental Psychology: Human Perception & Performance*, 18(1), 106–120.
- Shipley, T. F., & Kellman, P. J. (1992b). Strength of visual interpolation depends on the ratio of physically specified to total edge length. *Perception & Psychophysics*, 52(1), 97–106.
- Singh, M., & Hoffman, D. D. (1999). Completing visual contours: The relationship between reliability and minimizing inflections. *Perception & Psychophysics*, 61(5), 943–951.
- Spehar, B., & Clifford, C. (2000). When does illusory contour formation depend on contrast polarity. *Vision Research*, 13(18), 1915–1919.
- Sugita, Y. (1999). Grouping of image fragments in primary visual cortex. *Nature*, 401, 269–272.
- Ullman, S. (1976). Filling in the gaps: The shape of subjective contours and a model for their generation. *Biological Cybernetics*, 25, 1–6.
- Victor, J., & Conte, M. (2000). Illusory contour strength does not depend on the dynamics or relative phase of the inducers. *Vision Research*, 40, 3475–3483.
- von der Heydt, R., Peterhans, E., & Baumgartner, G. (1984). Illusory contours and cortical neuron responses. *Science*, 224, 1260–1262.
- Wertheimer, M. (1923). Laws of organization in perceptual forms. First published as Untersuchungen zur Lehre von der Gestalt II, in *Psychologische Forschung*, 4, 301–350 [Translation published in Ellis, W. (1938). A source book of Gestalt psychology (pp. 71–88). Kegan, Paul, Trench & Teubner, London].
- Williams, L. R., & Jacobs, D. W. (1997). Perceptual completion of occluded surfaces. *Neural Computation*, 9(4), 837–858.
- Wurtz, R., & Lourens, T. (2000). Corner detection in color images through a multiscale combination of end-stopped cortical cells. *Image and Vision Computing*, 18, 531–541.
- Yazdanbakhsh, A., & Grossberg, S. (2004). Fast synchronization of perceptual grouping in laminar visual cortical circuits. *Neural Networks*, 17, 707–718.
- Yen, S. C., & Finkel, L. H. (1998). Extraction of perceptually salient contours by striate cortical networks. *Vision Research*, 38, 719–741.
- Yin, C., Kellman, P. J., & Shipley, T. F. (1997). Surface completion complements boundary interpolation in the visual integration of partly occluded objects. *Perception*, 26, 1459–1479.
- Zemel, R., Behrmann, M., Mozer, M., & Bavelier, D. (2002). Experience-dependent perceptual grouping and object-based attention. *Journal of Experimental Psychology: Human Perception & Performance*, 28(1), 202–217.