Journal of Experimental Psychology:
Human Perception and Performance
2005, Vol. 31, No. 3, 558–583

0096-1523/05/$12.00 DOI: 10.1037/0096-1523.31.3.558

3-D Interpolation in Object Perception: Evidence From an Objective Performance Paradigm

Philip J. Kellman and Patrick Garrigan
University of California, Los Angeles

Thomas F. Shipley
Temple University

Carol Yin and Liana Machado
University of California, Los Angeles

Object perception requires interpolation processes that connect visible regions despite spatial gaps. Some research has suggested that interpolation may be a 3-D process, but objective performance data and evidence about the conditions leading to interpolation are needed. The authors developed an objective performance paradigm for testing 3-D interpolation and tested a new theory of 3-D contour interpolation, termed *3-D relatability*. The theory indicates for a given edge which orientations and positions of other edges in space may be connected to it by interpolation. Results of 5 experiments showed that processing of orientation relations in 3-D relatable displays was superior to processing in 3-D nonrelatable displays and that these effects depended on object formation. 3-D interpolation and 3-D relatabilty are discussed in terms of their implications for computational and neural models of object perception, which have typically been based on 2-D-orientation-sensitive units.

Object perception is basic to thought and behavior. Its function is to provide representations of coherent, connected entities in the world. When it is accurate, object perception tells perceivers which parts of the world will cohere when acted upon and where things will separate. Many properties of objects are important: shape, size, composition, function, and so on. Yet the most basic property is the one that makes objects so important in representations of reality: Perceived objects correspond to units in the physical world.

Humans' most powerful sense in perceiving objects is vision. Even at a distance, a person can see where the world is divided into parts. Object perception in ordinary circumstances is fast and accurate, yet its ease and success conceal its complexity. Although research in recent years has facilitated understanding of aspects of visual segmentation and grouping processes that lead to perceived objects (for a review, see Shipley & Kellman, 2001), many aspects remain poorly understood.

A particularly important and difficult problem in understanding object perception is how the visual system copes with fragmentation in the input. Most objects in ordinary scenes are partly occluded; yet observers are able to obtain representations of whole objects under most circumstances. A great deal of research suggests that the visual system uses interpolation processes to connect visible contours and surfaces across spatial gaps in the optical projections of objects (Heitger, von der Heydt, Peterhans, Rosenthaler, & Kubler, 1998; Kanizsa, 1979; Kellman, Guttman, & Wickens, 2001; Kellman & Shipley, 1991; Michotte, Thines, & Crabbe, 1964; Petry & Meyer, 1987; Thornber & Williams, 2000).

Although representations of objects, or at least some aspects of objects, include 3-D information (e.g., Liu, Knill, & Kersten, 1995), most research on visual interpolation has focused on 2-D displays. Some research has begun to address the third dimension in visual object formation (e.g., Carman & Welch, 1992; Hakkinen, Liinasuo, Kojo, & Nyman, 1998; Heider, Spillman, & Peterhans, 2002; Kellman & Shipley, 1991), but there has been little in the way of objective data or formal theory regarding 3-D interpolation. Recently, Kellman, Garrigan, and Shipley (in press) proposed a systematic theoretical account of 3-D object formation. The theory of 3-D relatability specifies the conditions governing interpolation of contours across gaps in 3-D space. In the research reported in the present article, we aimed to develop an objective performance paradigm to test 3-D interpolation as well as some specific predictions of 3-D relatability regarding when interpolation does and does not occur. Specifically, we asked, using an objective performance method, whether object interpolation processes use as inputs the 3-D positions and orientations of edges and surfaces. We also sought to show that object interpolation processes produce as outputs representations of contours and surfaces that extend through all three spatial dimensions. We studied both of these questions through their effects on perceptual processing. The results provide clear evidence that grouping and segmentation are 3-D processes and that their geometry is described by 3-D relatability.

Not much is known about 3-D interpolation from prior research. A number of investigators have explored important aspects of 3-D segmentation and grouping, but most often these explorations have involved depth stratification of two or more frontoparallel planes.

Philip J. Kellman, Patrick Garrigan, Carol Yin, and Liana Machado, Department of Psychology, University of California, Los Angeles; Thomas F. Shipley, Department of Psychology, Temple University.

Liana Machado is now at the Department of Psychology, University of Otago, Dunedin, New Zealand.

This research was supported in part by National Eye Institute Grant RO1 EY13518 to Philip J. Kellman and Thomas F. Shipley. We thank John Foley and Carlo Fantoni for helpful discussions and Jennifer Vanyo for general assistance.

Correspondence concerning this article should be addressed to Philip J. Kellman, Department of Psychology, University of California, Los Angeles, Los Angeles, CA 90095-1563. E-mail: kellman@cognet.ucla.edu

(For a review of research relevant to 3-D interpolation, see Kellman et al., in press). In the following sections, we note prior work that directly involves 3-D interpolation—that is, cases in which interpolation depends on the 3-D positions and relations of physically specified edges.

Illusory Contours

Kellman and Shipley (1991) suggested that contour interpolation processes use the positions and orientations of inducing edges in 3-D space. They produced a stereoscopic demonstration in which interpolated boundaries smoothly varying in three dimensions were seen. A similar figure is shown in Figure 1. The object has both a 3-D illusory contour and a 3-D partially occluded contour, consistent with the idea that these phenomena have a common interpolation process (see Kellman, Yin, & Shipley, 1998; Shipley & Kellman, 1992a). Differences in their appearances arise from differences in the completed contour's depth relative to other surfaces (Kellman et al., in press; Kellman & Shipley, 1991). This account is compatible with the reversal seen when the views for the left and right eyes are switched: The portion of the contour that previously appeared as illusory now appears occluded, and vice versa.

Experimental investigations of similar displays have found that a variety of shape classes may be seen in stereoscopic illusory figure displays (Carman & Welch, 1992) and that there is systematic underperception of curvature (Vreven & Welch, 2001). In both cases, the authors argued that the phenomenon depended on interpolated contours and surfaces. These interpretations are plausible, but the research was based exclusively on subjective report. Because factors other than perceptual representations can influence such reports, it would be helpful to have data from an objective paradigm to establish the existence of 3-D interpolation and understand its mechanisms.

In addition, previous research has not sought to define or test the conditions hypothesized by Kellman et al. (in press) to produce 3-D interpolation, although one study reported preliminary findings in this area (Kellman, Machado, Shipley, & Li, 1996). The research reported in the present article addressed the spatial relations among 3-D contours required to support 3-D completion.

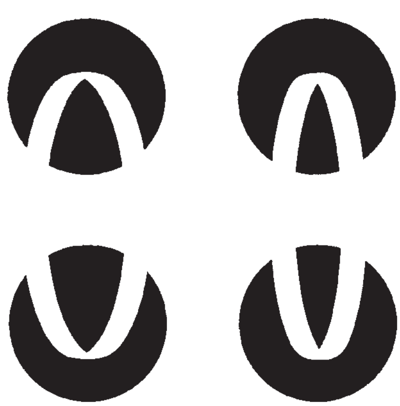

Figure 1. 3-D illusory contours. When free-fused, this stereo pair produces the perception of a ring turned out of the picture plane. The figure has a pair of occluded and illusory contours on its left and right sides. Switching the left- and right-eye images also switches which pair of contours appears occluded and which pair appears illusory. From "A Theory of Visual Interpolation in Object Perception," by P. J. Kellman and T. F. Shipley, 1991, *Cognitive Psychology, 23,* p. 182. Copyright 1991 by Elsevier. Adapted with permission.

Object Completion and Contour Relatability

Both theoretical and experimental work have considered how the visual system identifies which parts of a visual scene are parts of the same object in the world. The problem, put simply, is how to identify the parts of an array that belong together when adjacent areas in the projection may come from different objects and a single object may project to many spatially separated regions. Significant progress has been made on the initiating conditions and geometric relationships of object completion in two dimensions. Here, we briefly review some of this work because it will be helpful in understanding completion in three dimensions.

Object completion appears to require complementary contour- and surface-completion processes (Grossberg & Mingolla, 1985; Kellman & Shipley, 1991). Contour completion is primary, in that surface properties can spread and are confined within real and interpolated boundaries as well as within extrapolated boundaries (Yin, Kellman, & Shipley, 1997, 2000). Both processes are likely relevant to 3-D completion, but we focus here on contour-completion processes because they lead in establishing both unity and shape.

The initiating conditions for contour completion are contour junctions—corners in the visible boundaries of a surface. Formally, these are discontinuities in the first derivatives (tangent discontinuities or first-order discontinuities) of the contour—points where the slope of the contour is undefined. Shipley and Kellman (1990) noted that interpolated contours in general extended between such points and that removing them—by rounding corners, for example—eliminated, or markedly reduced, contour completion. Heitger et al. (1998) referred to these points as *key points* and proposed a neurally plausible model for their extraction from images.

The presence of tangent discontinuities (TDs) is necessary but not sufficient for contour completion. TDs are ordinarily present when an edge is occluded, but they can also occur in the natural boundaries of objects. Identification of when contour completion is necessary requires other criteria. Kellman and Shipley (1991) have proposed that the visual system uses specific geometric relations among contours leading into TDs to interpolate connections. These geometric relations have been formally characterized in terms of *contour relatability* (Kellman & Shipley, 1991; Singh & Hoffman, 1999a; cf. Geisler, Perry, Super, & Gallogly, 2001). Two contours are relatable if they meet three constraints: smoothness (the interpolated contour can be differentiated at least once), monotonicity (interpolated contours can bend in only one direction), and (roughly) a 90° limit (interpolated contours can bend up to 90°). Figure 2 shows a construction that is useful in defining contour relatability. Formally, if E_1 and E_2 are surface edges, and R and r are perpendicular to these edges at points of TD (assigned such that $R \geq r$), then E_1 and E_2 are relatable if and only if

$$0 \leq R \cos\theta \leq r,\ R \geq r.$$

The second part of this inequality expresses the conditions required for a smooth, monotonic curve connecting E_1 and E_2 that

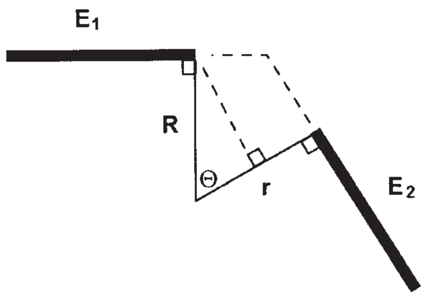

Figure 2. Geometric definition of 2-D relatability. Two edges (E_1 and E_2) are relatable if and only if the two perpendiculars R and r (chosen so that $R \geq r$) extending from the ends (tangent discontinuities) of E_1 and E_2 meet and the angle between E_1 and E_2, θ, is bounded by the relationship $0 \leq R \cos\theta \leq r$.

agrees with their orientations at their points of TD. The first part expresses the 90° limit, in that $\cos\theta$ becomes negative when θ exceeds 90°.

Limits of Relatability and Quantitative Variation

The formal criterion given above specifies the limits of relatability—that is, the range of positions and orientations of edges that support interpolation. Limits are important: A crucial question in object formation is whether two visible fragments are or are not parts of the same object. This determination probably governs whether certain kinds of further processing will occur. Seeing visual patches as parts of a single object, or not, influences attention and subsequent processing (Baylis & Driver, 1993; Zemel, Behrmann, Mozer, & Bavelier, 2002).

Not all relatable contours connect with equal strength (Banton & Levi, 1992; Field, Hayes, & Hess, 1993; Kellman & Shipley, 1991; Shipley & Kellman, 1992a, 1992b; Singh & Hoffman, 1999b). Factors known to influence the strength of interpolation include support ratio (Banton & Levi, 1992; Shipley & Kellman, 1992b), the turning angle of curved interpolation (Field et al., 1993), and misalignment (Shipley & Kellman, 1992a). Figure 3 illustrates changes in perceived unity as the angle between edges deviates from collinearity (up to an approximate limit at 90°). In a contour-detection paradigm, Field et al. (1993) found that a path of Gabor elements, defined by the position and orientation relationships among the elements, was not detectable if the local orientation difference between adjacent elements was 90°. (For a quantitative proposal of the relationship between relatability and angular difference, see Singh and Hoffman, 1999a.)

Relatability and Good Continuation

Relatability is sometimes equated with the Gestalt principle of *good continuation* (Wertheimer, 1923/1938). Although these two concepts are related, there are important distinctions between them (Kellman, Garrigan, Kalar, & Shipley, 2003). Good continuation and relatability are both based on an assumption about the "smoothness" of contours in the world (Marr, 1982). However, as stated by Wertheimer (1923/1938), good continuation deals with the segmentation of visible parts, whereas relatability is primarily concerned with the connecting of contours through areas in which the connection is not visible. Clear differences are evident in the influences on these two tasks. Segmenting of visible regions is not influenced by support ratio, and what counts as a unitary contour segment does not appear to be limited by whether the contour (if smooth) turns through more than 90°. Experiments (Kellman et al., 2003) have indicated that TDs are crucial in the segmentation of continuous contours into perceived parts. Segments that are zero-order and first-order continuous (i.e., no TDs) are perceived as unitary. Relatability applies under different conditions—specifically, situations in which there are zero-order discontinuities in the stimulus (gaps requiring interpolation). Contour interpolation is triggered by TDs (first-order discontinuities), but rather than leading to segmentation, these lead to connection with other contours. Relatability also incorporates several constraints not present in good continuation (including monotonicity and the 90° constraint). This analysis applies to Wertheimer's notion of good continuation. One difficulty with the notion of good continuation is that it has been used in diverse and often vague ways. In contour interpolation specifically, good continuation has been used to refer to linear extensions under occlusion (Fantoni & Gerbino, 2003). This notion also clearly differs from relatability, which can form smooth, curved connections between straight inducing edges. What relatability and good continuation have in common is that both involve smoothness constraints. More than one smoothness principle is needed to do different kinds of perceptual work in visual segmentation and grouping, however (Kellman et al., 2003).

3-D Relatability

What geometric relations describe the linking of edges oriented in three dimensions into a 3-D object? The 3-D case requires more than 2-D relations. Indeed, the 2-D relations that produce interpolation appear to be a special case in a more general formulation of relatability. Kellman et al. (in press) discussed a detailed discussion of 3-D relatability and its relation to the 2-D case, and they provided a detailed formulation of 3-D relatability. Here, we summarize the main ideas of 3-D relatability.

3-D contour relatability is a logical extension and generalization of the principles of 2-D contour relatability. The key constraints on interpolated contours in two dimensions may be applied to 3-D. In particular, Kellman et al. (in press) have theorized that interpolated

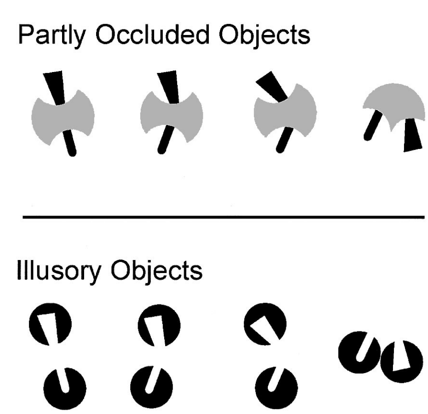

Figure 3. Systematic variation of relatability with deviation of edge orientation from colinearity for partly occluded and illusory objects.

contours in three dimensions must begin and end at TDs, must be smooth and monotonic, and must match the orientation of the visible contours that they continue at their TDs. If these criteria do apply to 3-D cases, then 3-D relatability seems likely to subsume 2-D relatability. Contour interpolation in frontoparallel planes, although it has dominated research, may actually be a special case of 3-D contour interpolation.

One way to express the requirements of 3-D relatability is as follows: Relatable edges must, within some threshold, be coplanar, but not necessarily coplanar in the frontoparallel plane, and within that plane, the edges must meet the 2-D constraints on relatability (for details, see Kellman et al., in press). The description involving 3-D coplanarity perhaps describes intuitively which edges are 3-D relatable, but it does not suggest how contour relatability may be realized in a neural architecture. A more suggestive formulation may build on an analogy with neural-style models for 2-D interpolation. Neural-style implementations of 2-D interpolation depend on relations of activation within a network of orientation-sensitive neural units, such as those known to exist in early cortical visual areas. This description seems appropriate for 2-D relatability, but it seems less so for contour interpolation in three dimensions. The reason is that 2-D models take as their inputs edge orientations and positions as these are encoded on the retina. If interpolation depends on 3-D positions, orientations, and relations, these models will not be adequate.

Suppose, however, that interpolation derives from interactions of oriented units that encode 3-D positions and orientations. Formally, these interactions define, for a given edge and any arbitrary point, the range of orientations at that point that fall within the limits of relatability (Kellman et al., in press). In a Cartesian coordinate system, we specify two angles, Θ and ϕ. Θ is the angle in the x–y plane, with 0 defined as the orientation parallel to the x-axis. ϕ is the angle in the x–z plane, with 0 also defined as the orientation parallel to the x-axis. For convenience, we place one edge so that its tip ends at the origin of the coordinate system (0, 0, 0), with an orientation $\Theta = \phi = 0$. For an edge terminating at any other point (x, y, z), we define the range of possible orientations (Θ, ϕ) for 3-D relatable edges terminating at that point. These are given by

$$\tan^{-1}\left(\frac{y}{x}\right) \leq \Theta \leq \frac{\pi}{2}$$

and

$$\tan^{-1}\left(\frac{z}{x}\right) \leq \varphi \leq \frac{\pi}{2}.$$

In these equations, the lower bounds express the absolute orientation difference ($-180°$) between the reference edge (terminating at the origin) and an edge ending at the arbitrary point oriented so that its linear extension intersects the tip of the reference edge. The upper bounds incorporate in three dimensions the 90° constraint.

Figure 4 shows some examples of edge positions and orientations that meet these criteria. We might call the set of relatable orientation and position combinations the 3-D *relatability field*. As described here, 3-D relatability specifies limits; it does not define quantitative variation within those limits. As a vector field, in other words, it assigns the value 0 everywhere outside the limits of relatability and positive values within the limits; but it does not yet assign variable strengths of interpolation that depend on position and orientation of edges within the limits of relatability.

Figure 4. 3-D relatable edges. The display is a stereo pair that can be free-fused by crossing the eyes. For one surface edge in three dimensions the set of relatable edges of any given orientation lie not on a plane (as they do in two dimensions), but rather fill a cone-shaped volume.

We would expect that for a given edge E, there would be variation in the strength of interpolation within the set of edges relatable to it. This variation would depend on at least two properties: (a) The distance over which interpolation must occur and (b) the angle between E and the members of the set of relatable edges. In three dimensions, as in two dimensions, interpolation should be most robust when edges are collinear and less so as the edges deviate from collinearity, with essentially no interpolation occurring when two edges meet at any angle more acute than 90°. Interpolation over short distances would be stronger than interpolation over longer distances.

The idea that interpolation depends on truly 3-D interactions, as suggested by this theory, would have a number of implications for models of object formation. Among these would be the idea that neural units encoding the relevant properties must exist and interact. After considering experimental tests of 3-D relatability, we return to this issue in the General Discussion. In this article, we report experiments on 3-D relatability in 3-D contour interpolation. Our hypotheses are that interpolation depends on positions and orientations of edges in 3-D space and that it is governed by the geometry of 3-D relatability.

To test these hypotheses, we obtained evidence regarding variations in strength of interpolation based on different amounts of depth misalignment, and we tested depth interpolation at several different 3-D angles of intersection. We expected that the hypothesis of 3-D relatability would apply to both occluded and illusory contours, but in the research presented here, we used illusory-contour stimuli. This choice was motivated by several considerations, the most important of which was the need to remove TDs from the stimuli in Experiment 4.

An Objective Paradigm for Testing 3-D Object Completion

Many of the phenomena of perceptual segmentation and grouping were initially conveyed through examples and demonstrations (e.g., Kanizsa, 1979; Michotte et al., 1964; Wertheimer, 1923/1938). More systematic perceptual report procedures, such as magnitude estimation, have also revealed important information about these phenomena and the visual processes that produce them (e.g., Day & Kasperczyk, 1983; Dumais & Bradley, 1976; Shipley & Kellman, 1992b).

Both demonstrations and perceptual report measures are useful, but they are limited by demand characteristics and by the influences of cognitive strategies in addition to the perceptual processes that they aim to assess. In contour interpolation in particular, it has been argued that experience with illusory contours leads to greater willingness to report them (Rock & Anson, 1979; Wallach & Slaughter, 1988).

There have been very few empirical studies of 3-D illusory contours, and none that have used objective measures. Carman and Welch (1992) asked subjects to report the shape seen in several 3-D illusory figures. Subjects spontaneously responded with the correct shape of the illusory figure, chosen from among four shape classes (planar, parabolic, elliptic, and hyperbolic). Each of these displays was a modified Kanizsa square, with the depth of the corners specified by stereodisparity. These results suggest that 3-D illusory surfaces can convey reliable shape information; however, they are limited by the difficulties associated with scoring free reports and by potential demand characteristics. Two of the 5 subjects were nonnaive, and it has been shown that awareness of what illusory surface is consistent with the shape arrangement of the inducing elements can bias performance (Wallach & Slaughter, 1988). In general it can be difficult in perceptual report studies of illusory contours to distinguish perception from a subjects' awareness that a certain shape could fit within or connect the inducing elements.

Because of these sorts of issues, objective performance methods have in recent years become standard in investigations of perceptual organization (e.g., Guttman, Sekuler, & Kellman, 2003; Ringach & Shapley, 1996; Sekuler, Palmer, & Flynn, 1994). Yet devising good objective methods is challenging because when perceptual organization is the issue, there is no objectively correct answer (e.g., about whether there is an illusory contour in a display). A solution is to devise an objective performance task on which performance will vary depending on the observer's perceptual organization of the input. If certain conditions are theorized to produce a certain organization, it can be predicted that the observer will consistently perform better (or worse) on some task, in which having that organization should help (or hurt).

In the present study, we devised a task analogous to one invented earlier to study 2-D illusory and occluded contours by Ringach and Shapley (1996). Their *fat–thin* task involves identifying the shape of a modified Kanizsa square whose inducing elements have been rotated to give the impression of a square that bulges horizontally and is short vertically (fat) or are in the opposite configuration (thin). Evidence suggests that performance on this task is facilitated by contour interpolation. Although judging the relative rotation of the inducing elements is possible whether or not illusory contours form, Ringach and Shapley found that performance under conditions thought to induce interpolation was better than performance in several control groups. Recent research using an image-classification paradigm (Beard & Ahumada, 1998) to detect which parts of displays are most influential in determining subjects' responses has provided further evidence that interpolated contours are used in this task (Gold, Murray, Bennett, & Sekuler, 2000).

For a number of reasons, it would be difficult to apply the fat-thin task directly to 3-D interpolation. One reason is that perspective effects involved with slanting surfaces in depth would be confounded with projective fatness or thinness of the displays. Another concern is that the fat–thin task can technically be performed by an ideal observer viewing any one of the inducing elements (although evidence indicates that subjects do in general use interpolated edges to perform the task). Because the phenomena of 3-D interpolation have been less studied and accepted than have those in the 2-D case, we aimed to construct a task that forced the subjects to use relations among elements.

Figure 5 shows the stimuli we used to study 3-D interpolation. Each panel shows a stereo pair that may be free-fused by crossing the eyes. When the stereo images are cross-fused, an observer should see two white areas on top of the black, planar "tabs." The white areas should appear in front of the black tabs and be arranged in a manner consistent with the side views shown beneath each display.

These stimuli have two important, orthogonal properties. First, as defined by the columns of Figure 5, the displays on the left each have two planes with 3-D relatable geometry. These planes appear to be connected by illusory contours. On the right, similar displays are shown, but here the geometry does not conform to 3-D relatability. These planes do not appear to be connected, and there is no percept of illusory contours. The second important property of these stimuli is independent of the relatable–nonrelatable criterion and is used to define our objective measure. Displays in the top row have two planes that are converging (intersecting). In the bottom row, each display has two planes that are parallel—either coplanar or not intersecting. Subjects in our experiment were instructed to quickly classify stimuli like these as either converging or parallel. This design allowed us to look for performance differences between 3-D relatable and nonrelatable stimuli on a task that required consideration of the stimulus as a whole and was therefore likely to be easier when the parts of the stimulus were perceptually connected.

Experiment 1

Using the parallel–converging method described above, Experiment 1 tested the hypothesis that 3-D completion, based on 3-D relatability, would produce accuracy and/or speed advantages in the parallel–converging task.

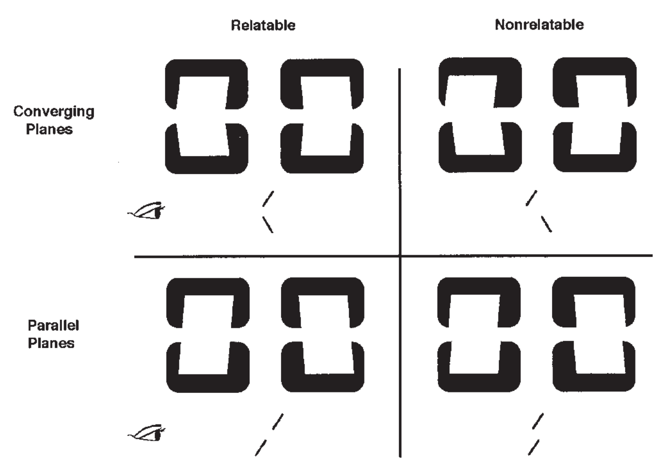

Figure 5. An objective 3-D classification task for illusory contours. Each display is a stereo pair that can be free-fused by crossing the eyes. The stimuli fall into two categories: intersecting (converging) planes (top) and parallel planes (bottom). In each quadrant, the upper image is a stereo pair of the two illusory planes oriented in depth, and the lower image is a side view of the same planes.

Method

Subjects. The subjects were 17 University of California, Los Angeles (UCLA) undergraduate students who received class credit for participation. An additional 3 subjects were not included in the final data set due to failure to meet a threshold criterion for performance on the experimental task (average d' [across conditions] $> .50$). All subjects had normal or corrected-to-normal vision. Each subject passed a simple screening for stereoscopic vision. A random-dot stereogram was presented on the monitor, and subjects had to report to the experimenter what form was present, floating in front of the background. If a subject gave the correct answer (a question mark), he or she was allowed to participate in the experiment. This test is a fairly demanding test of stereovision, because it is known that depth perception takes longer in random-dot stereograms than in comparable normal stereograms (Howard & Rogers, 1995).

Apparatus. The experiment was run on a Macintosh Power PC computer with an E-Machines TX-21 monitor at a resolution of 832 × 634 pixels. Each pixel was 0.435 mm in height and width and subtended 1.40 arcmin at subjects' viewing distance of 107 cm.

Subjects were positioned in a headrest at a viewing distance of 107 cm from the monitor. The headrest centered the subject in front of the displays and prohibited head movements, ensuring that disparity information was received from the correct viewing position. Stereoscopic images were presented using CrystalEyes LCD shutter glasses synchronized to the computer monitor such that the shutter over each eye was opened electronically while the appropriate image for that eye was displayed on the monitor.

Stimuli. Examples of the stimuli used in this experiment are shown in Figure 5. Each display consisted of two illusory-figure-inducing elements (the black forms with white cut-out sections). In the experiment, these inducing forms were red in a black surround. (Use of the red phosphor only in the CRT monitor minimized crosstalk between the two eyes' views.) Differences in the black cutout regions between the two eyes were used to create two black tabs (one in the upper and one in the lower part of each display). The tabs could take on a variety of slants and positions in depth, but they always appeared in front of the red forms. The inducing elements subtended visual angles of 3.35° horizontally and 2.47° vertically and were separated by 0.72°. The outside corners of the inducing elements were rounded slightly (radius of curvature = 12.6 arcmin) to avoid their producing illusory contours along their outside vertical edges. (The importance of TDs in illusory-contour formation is discussed in Experiment 4.) The cutout sections (tabs) were 2.21° wide × 1.44° high, and they were separated vertically by 0.72°. These characteristics produced a contour interpolation support ratio (Shipley & Kellman, 1992b) of .8, which was held approximately constant in Experiments 1–5. (Support ratio of interpolated edges is defined as the ratio of the length of the physically specified parts of an edge to the total edge length—i.e., physically specified extent plus gap.)

In our remaining descriptions of the stimuli, we give the characteristics and 3-D positions of the virtual objects as they were specified to the observers during stereoscopic viewing. Where useful, we also refer to specific image manipulations that were used to obtain these virtual objects.

Positions. Tabs in the displays could appear at three virtual distances from the observer, as specified by binocular disparity: 91.0, 95.5, and 100.0 cm. For convenience, we refer to these as Positions 1, 2, and 3, with Position 1 being closest to the observer.

Slants. By *slant* we mean rotation around a horizontal axis perpendicular to the viewer's line of sight. A frontoparallel tab is said to have slant of 0. We refer to cases in which the top edge is further from the viewer than the bottom edge as *positive slants*; *negative slants* refer to cases in which the bottom edge is further from the viewer than the top edge. Tabs could have any of eight slants in the experiment: approximately 14°, 26°, 46°, and 64°, at positive or negative slants for each. Exact slant values varied slightly depending on the tabs' positions (i.e., observer-relative distance); these variations ranged from about ±1.5° at the smallest slant to about ±3.5° at the largest slant.

Parallel and converging planes. All tabs were presented at a base disparity so that they always appeared entirely in front of the two red inducing elements. Each tab in each pair was slanted in depth as if rotated around a horizontal axis passing through the vertical midpoint of its physically defined region. Slant magnitude was specified solely through stereodisparity. (In other words, vertical and horizontal extents of the tabs were held constant, and stereodisparity was created by shearing the vertical edges in opposite directions in the two eyes.) Relations between the planes of the two tabs in each display could be either parallel or converging. For parallel displays, the top and bottom tabs were in parallel planes. Another way of describing parallel displays is that the two defined tabs always had identical directions and magnitudes of slant. In converging displays, the top and bottom tabs lay in intersecting planes—that is, the two tabs were slanted in opposite directions but had the same slant magnitudes.

Relatable displays: Definition. Subsets of the parallel and converging displays met the criteria of 3-D relatability as given above. In all cases, relatable displays were positioned so that their vertically oriented contours were coplanar (but not in a frontoparallel plane) and could be connected by a smooth, monotonic, interpolated contour that matched the real, physically defined contours' orientations at their endpoints. Preliminary observations suggested that this relationship of 3-D relatability between the upper and lower black tabs led to formation of illusory contours and surface connections between them.

Relatable displays: Construction. Subsets of both the parallel and converging displays were relatable. For parallel displays, this meant that the two slanted tabs, not just their vertical bounding contours, were coplanar. Thus, the two tabs were at different distances from the observer. If perceptual completion occurred between the two tabs, the appearance of these displays would be a planar surface connecting the two tabs, slanted in depth. For the subset of converging displays that were relatable, the tabs were equidistant from the observer. Relatable displays were placed equally often at Positions 1, 2 and 3. This manipulation served to equate the average distance from the observer of relatable and nonrelatable tabs across the whole experiment.

Nonrelatable displays. Another set of displays was nonrelatable relative to the definition given above. From each of the displays defined as relatable, nonrelatable displays were created by shifting one of the tabs in depth, using stereodisparity. For example, a shift further in depth was achieved by moving all points of the left eye's view leftward by some amount and all points in the right eye's view rightward by the same amount. In the converging condition, this shift necessarily resulted in the two black tabs being at different observer-relative depths. In the parallel condition, the shift could bring the tabs closer in depth or further apart in depth, depending on the slant, whether the top or bottom tab was shifted, and the direction of shift (see Figure 6). Two levels of shift were used to create two different magnitudes of departure from relatability: These consisted of about a 4.5-cm and 9.0-cm shift in the virtual displays, with the average viewing distance being 95 cm. The nonrelatable displays equally often involved shifts of the top and bottom tabs for each angular slant in both the parallel and converging conditions. Nonrelatable displays having a 4.5-cm shift appeared equally often at Positions 1 and 2 and at Positions 2 and 3. Nonrelatable displays having a 9.0-cm shift always occupied Positions 1 and 3.

The design of the stimuli met several objectives of experimental control. First, to rule out certain monocular bases of response in the experimental task, perspective cues were held constant. The overall height and width of the images was held constant for all displays. (This meant that the dimensions of the specified virtual objects varied somewhat with slant and shift.) Second, we decoupled, as much as possible, the information for depth shift and for slant. This was accomplished by varying the relative positions of the two eyes' images to produce shift and by using shear to produce slant. The virtue of this arrangement was, again, avoidance of perspective cues

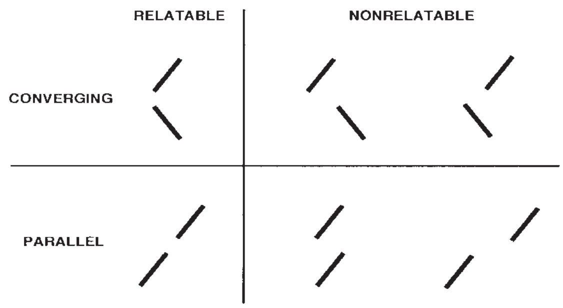

Figure 6. Depth relations in the experimental displays. Nonrelatable converging stimuli always spanned a larger depth range than relatable converging stimuli. Nonrelatable parallel stimuli could span a larger or smaller depth range than relatable parallel stimuli, depending on the direction of shift and the slant of the inducing elements.

that would correlate with, and might provide information about, depth shift or slant. Accordingly, however, the specified slant in the virtual object also varied slightly with depth shift. These differences were small (<3.5° in the most extreme case), especially relative to subject variability in slant perception (see Experiment 5). The strategy behind these design choices was to minimize proximal stimulus differences in each eye across conditions, with the consequence that the virtual objects varied somewhat. One issue raised by these choices is that slant may have been underperceived in these displays. Because horizontal size ratio (between the left and right eyes' views) was always 1, the widths of the horizontal edges of the virtual objects that were further from the observer were slightly wider than were parts closer to the observer. This characteristic of the stimuli did not create a depth cue conflict per se but did create the property that each object had a rectangular cyclopean projection. This property was largely irrelevant for the purposes of Experiment 1, but it has relevance to the exact amount of slant perceived in the displays, a topic we consider in Experiment 5.

Design. The experimental task consisted of 384 trials. On each trial, a subject saw a display and made a forced-choice judgment of whether the two tabs in the display were parallel or converging. Parallel and converging displays each constituted half of the trials. Orthogonal to this classification, one third of the displays were relatable, one third were shifted in depth by 4.5 cm, and one third were shifted by 9.0 cm. Given two relative orientations (parallel vs. converging), three shift values, and four slants, there were 24 basic displays. Each relatable display type was presented 24 times, and each nonrelatable display type was presented 12 times. Because there were two levels of nonrelatable displays (two nonzero shift values), this arrangement produced equal numbers of trials for relatable and nonrelatable displays in the experiment. Other factors, such as positive versus negative slant direction, top versus bottom shift, and the positions of relatable displays were counterbalanced across the trials involving each display.

Procedure. Each subject in the experiment was seated in a comfortable chair with his or her head position stabilized by an adjustable chinrest. Subjects were told they were to participate in an experiment involving depth perception. They were told that they would be wearing 3-D goggles enabling them to see depth in displays shown on the monitor. The experimenter explained that each display would have an upper and a lower tab and that these tabs could be in either converging (intersecting) or parallel planes. A physical model was presented to convey these ideas. It consisted of two thin cardboard tabs attached to wires that extended through slits in two parallel walls (see Figure 7). The model allowed the two tabs to be placed in both parallel and converging relationships. Subjects were instructed that the two tabs would appear at varying depths, and several cases were shown. These included both relatable and nonrelatable examples, although no mention was made of these notions.

An ordinary computer keyboard was provided, and subjects were instructed to press *P* for parallel and *C* for converging tabs on each trial. They were also instructed to press the space bar to initiate each trial. Finally, they were told that in the actual experiment, a display would appear briefly, followed by a display of random dots (used as a mask).

Subjects were then fitted with the LCD shutter glasses and presented with 8 trials on which different types of displays from the experiment were shown. The subject was asked to tell the experimenter his or her classification of each display (parallel vs. converging). These trials allowed the experimenter to ascertain that the subjects understood the task and were able to perceive stereoscopic depth through the shutter glasses.

Subjects were then given a block of 57 practice trials during which they received feedback on their responses. (A low-pitched tone indicated that an answer was correct; three high-pitched beeps indicated an error.) Subjects were instructed to respond as quickly as possible while preserving accuracy.

At the completion of practice, the actual experiment began. No feedback was given during the experimental trials. An experimental session took under an hour and was divided into four parts, separated by short rest periods.

Dependent Measures and Data Analyses

To equate for the varying positions of tabs required in the shifted conditions, we averaged performance for relatable displays tested at the three different positions for the analyses. Likewise, the two possible locations of 4.5-cm shifted displays (Positions 1 and 2 or Positions 2 and 3) were used equally often in the experiment and averaged for the analyses. These procedures ensured that tabs in the relatable and in each shift condition appeared at the same average observer-relative distance across the experimental trials.

The main results of this experiment are shown in Figures 8 and 9. Figure 8 plots sensitivity (d') as a function of stereoscopically specified tab slant for relatable 4.5-cm shifted and 9.0-cm shifted displays. It appears that 3-D relatability exerted a strong effect on performance in the classification task. Subjects showed greater sensitivity for 3-D-relatable displays at all four of the slant values tested. As can be seen in Figure 9, speed was correlated with sensitivity: Faster responding occurred for 3-D-relatable displays.

These patterns were confirmed by the analyses. Sensitivity was

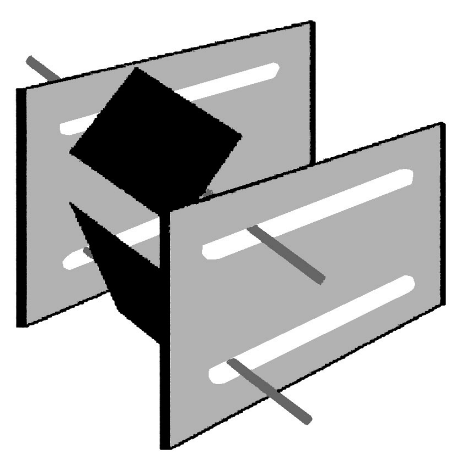

Figure 7. Demonstration apparatus.

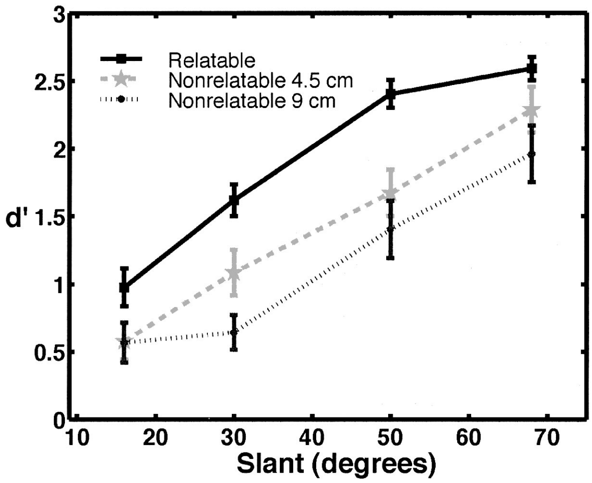

Figure 8. Sensitivity (d') as a function of slant in Experiment 1. Error bars represent plus or minus 1 standard error of the mean.

analyzed in a 3 (relatability: relatable, 4.5-cm shifted, 9.0-cm shifted) × 4 (slant) within-subject analysis of variance (ANOVA). There was a strong main effect of 3-D relatability, $F(1, 16) = 42.58$, $p < .001$. As expected, sensitivity increased as a function of increasing tab slant, indicated by a main effect of slant, $F(1, 16) = 109.91$, $p < .001$. There was no reliable Relatability × Slant interaction, $F(1, 16) = .013$, *ns*.

Individual comparisons indicated that the relatable displays differed reliably from both the 4.5-cm shifted, $t(16) = 6.53$, $p < .001$, and the 9.0-cm shifted displays, $t(16) = 8.12$, $p < .001$. The two shifted conditions also differed from each other, $t(16) = 2.69$, $p < .05$.

Response times (RTs) are shown for converging displays in Figure 9A and for parallel displays in Figure 9B. The plots indicate that subjects processed relatable displays more rapidly, on the order of 200 ms or about 12% on average. These trends were confirmed by the analyses. A 3 (relatability: relatable, 4.5-cm shifted, 9.0-cm shifted) × 2 (display type: parallel vs. converging) × 4 (slant) within-subject ANOVA was carried out. There was a main effect of relatability, $F(1, 16) = 22.89$, $p < .001$; a main effect of slant, $F(1, 16) = 39.28$, $p < .001$; but no main effect of display type, $F(1, 16) = 1.85$, *ns*. There was a reliable Relatability × Slant interaction, $F(1, 16) = 9.53$, $p < .01$, apparently a result of increasing differences among the three relatability conditions at increasing slants. There was also a reliable Display Type × Slant interaction, $F(1, 16) = 26.73$, $p < .001$, reflecting the more consistent linear decrease in RT with increasing slant for converging displays versus a more curvilinear pattern, with RTs differing only slightly for the two smallest slants, for parallel displays. There were no other reliable main effects or interactions.

Planned comparisons showed that all three relatability conditions differed from each other. 3-D-relatable displays differed from both nonrelatable 4.5-cm shifted displays, $t(16) = 4.79$, $p < .001$, and nonrelatable 9.0-cm shifted displays, $t(16) = 5.58$, $p < .001$. Nonrelatable 4.5- and 9.0-cm shifted displays differed modestly from each other, $t(16) = 2.23$, $p < .05$.

Discussion

The results of Experiment 1 support the idea that contour interpolation is a 3-D process. The classification of display tabs as parallel or converging was facilitated for 3-D relatable displays relative to displays in which 3-D relatability was disrupted by a depth shift of one tab. These results match up with the phenomenology of the displays in Figure 5: When a coplanar or otherwise smooth, monotonic connection can be made between the tabs, the visual system creates an interpolated surface connection between them. This surface is clearly bounded by illusory contours and bends smoothly in depth to link the two tabs. The most straightforward interpretation of Experiment 1 is that 3-D interpolation allowed more efficient performance in the classification task. Moreover, 3-D interpolation appears to depend on conditions of 3-D relatability—namely, that the input edges must be connectable by a smooth, monotonic curve in three dimensions. Disruption of 3-D relatability reduced or eliminated 3-D interpolation, shown by reduced sensitivity and speed in perceptual classification. The two

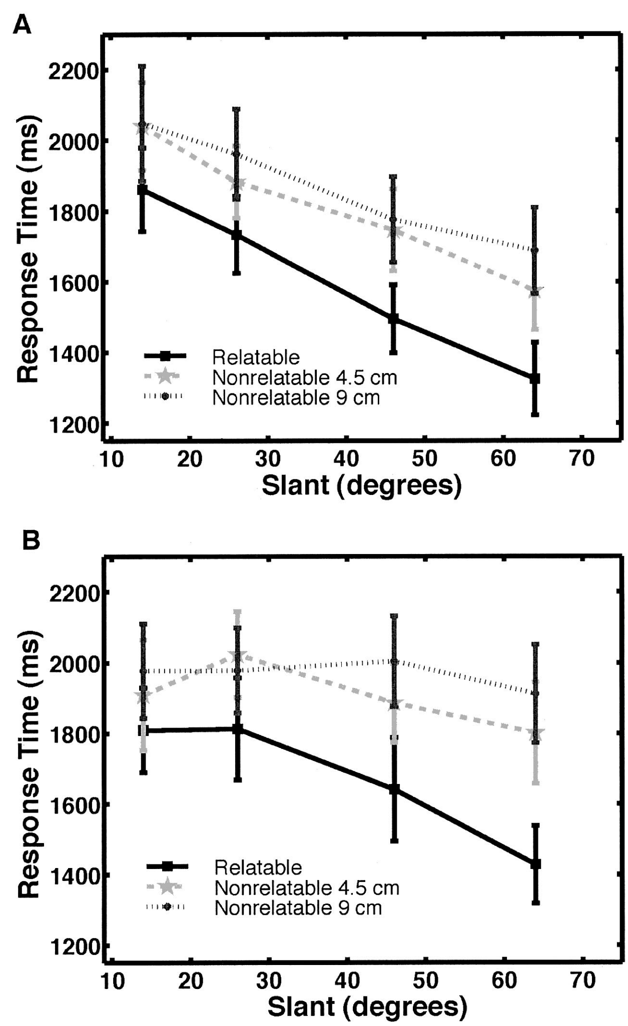

Figure 9. Response time as a function of slant for parallel (A) and converging (B) displays in Experiment 1. Error bars represent plus or minus 1 standard error of the mean.

levels of shift differed reliably, suggesting that strength of interpolation may not have decreased to zero in the smaller shift displays. Performance was worse for displays with the larger shift. Whether the larger shift completely eliminated interpolation cannot be decided from these data alone. Data from later experiments suggest that levels of performance in the larger shift condition are consistent with the absence of any interpolation effects.

In short, the results of Experiment 1 are consistent with positive answers to the two questions posed earlier. On the input side, object formation depends on the 3-D orientations and positions in 3-D space. On the output side, interpolation processes produce contours and surfaces that bend through all three spatial dimensions. Moreover, these results begin to define the specific requirements for 3-D interpolation. Specifically, interpolation depends on contour relationships satisfying a 3-D criterion of relatability—having smooth, monotonic connections that match the orientations of input edges at contour junctions.

Although plausible, these interpretations require further scrutiny. They depend in the first place on the validity of the method used. We hypothesized that our classification task would be sensitive to unit formation, in that perceptual organization of the tabs into a single object might permit faster and more accurate classification than would be found in conditions not leading to unit formation. Moreover, we hypothesized a particular set of geometric relationships—3-D relatability—that might define the relevant conditions for interpolation. Despite the clarity of the results, there is a certain amount of bootstrapping in both validating a method and identifying the relevant geometry of interpolation from the same set of data.

At least two alternative hypotheses must be considered. One concerns the method: Does it really access performance differences that are based on object completion? In other words, is the superiority of relatable over nonrelatable displays truly a completion effect? Perhaps some other aspect of the configurations of relatable and nonrelatable displays made classification of the former easier than classification of the latter. One possibility seems especially salient. In converging displays, relatable versions had two tabs at the same egocentric distance (distance from the observer). For the nonrelatable (depth-shifted) displays, the two tabs appeared at different egocentric distances. Perhaps our classification task, which depended on a comparison of the two tabs on each trial, was easier when these tabs were at the same egocentric distance.

Supposing the method was truly sensitive to object completion, there is also a question of whether the relevant completion is 3-D in nature. The displays in the three conditions were designed to differ as 3-D displays, but stereoscopic differences used to create the depth effects also induced monocular differences in displays across conditions. Perhaps 2-D completion processes, not 3-D ones, boosted performance in the relatable condition. We address this issue in Experiments 3 and 4. First, however, we take up the issue of whether the observed performance differences depend on object completion or on depth relations. (As is shown in Experiments 3 and 4, the most general version of this issue requires even further consideration.)

A completion effect? As discussed above, a straightforward interpretation of the results of Experiment 1 is that the effects of condition (3-D-relatable vs. shifted displays) derived from the fact that only the 3-D-relatable displays gave rise to object completion. It remains possible, however, that some characteristic of the relatable displays led to a performance advantage apart from any considerations of the visual system forming a connection between the two presented tabs.

There are a number of possibilities in this general class of alternative hypotheses. Aspects of Experiment 1 were designed, however, to test the most obvious of these—namely, that comparisons of tabs at equal egocentric distances might be easier than comparisons of tabs at different distances. Note that similarity or difference of egocentric distance of tab pairs is not perfectly correlated with the classification of displays as relatable or nonrelatable. For converging relatable displays, tabs were always at identical egocentric distances; corresponding nonrelatable converging displays always had tabs at different distances. For parallel relatable displays, however, the tabs were always at different egocentric distances. (Parallel relatable displays always contained two tabs that were coplanar—in a plane slanted toward or away from the observer.) To make nonrelatable displays in the parallel case, two kinds of shifts were used equally often. Half of these shifts moved a tab so that the two tabs became more similar in egocentric distance than they were in the relatable case; the other half moved a tab so that the two tabs differed even more in egocentric distance than they did in the relatable case.

This arrangement allowed us to perform an additional analysis within Experiment 1. This analysis compared subjects' responses for all of the parallel relatable displays in the experiment for which some corresponding nonrelatable displays produced smaller depth differences for the two tabs. If judging tabs at approximately the same egocentric distance was responsible for the performance advantage, we would predict that performance with this subset of nonrelatable displays would be superior to that with the corresponding relatable displays (in which tabs were more different in depth positions). In contrast, if the effect truly depended on object completion as given by 3-D relatability, we would expect the opposite result in this analysis.

This analysis showed that relatable displays (average depth difference = 2.98 cm) were more accurately classified than corresponding nonrelatable displays with lower average depth differences (average depth difference = 2.08 cm); mean proportions correct were .92 and .83 for relatable and nonrelatable displays, respectively, $t(16) = 1.81$, $p < .05$ (one-tailed). RTs were longer for the nonrelatable (M = 1,702.4 ms) than for relatable (M = 1,444.7 ms) displays, $t(16) = 4.82$, $p < .001$. This analysis of a subset of displays in Experiment 1 clearly indicates that 3-D relatability, not similarity of the tabs in depth, was the primary determinant of the performance difference between relatable and shifted displays.

Addressing the equal-depth hypothesis is an important step in verifying that the results of Experiment 1 indicate object-completion effects. Yet the idea that equal depth facilitates comparisons is only the most salient idea of this sort. There are others. For example, it might coincidentally be the case that the smoothness of the depth gradient between two tabs facilitates their comparison, apart from object completion. This sort of alternative explanation comes very close to the notion of relatability but, nevertheless, differs from it. It could be that the exact geometry of 3-D relatability is, just by coincidence, also an account of the best relations of pairs of surface fragments for making comparisons of the sort required in our task. In other words, relatability may identify a perceptually special set of spatial relations but not those involved in object completion. This issue, in part, motivated Experiments 3 and 4.

A depth effect? We turn now to the other type of alternative hypothesis. Supposing that the performance advantage with relatable displays depends on object completion, how do we know the effect necessarily involves 3-D completion? The issue can be seen in a consideration of some 2-D characteristics of the images used to create relatable and nonrelatable displays.

All displays in these experiments were stereo pairs, with disparate images shown to left and right eyes. In the relatable displays, the top and bottom tabs were always relatable in two dimensions—that is, in the frontoparallel plane of each image. This connection between 3-D and 2-D relatability is unavoidable. A theorem of projective geometry guarantees that (apart from degenerate cases) any smooth curve in 3-D space, when projected onto a plane, will produce a smooth curve in the 2-D projection (Gans, 1969). A small, additional requirement is needed for relatability—namely, that any smooth, monotonic curve in 3-D space projects to a smooth, monotonic curve in two dimensions. Although we do not know of an explicit proof of this claim, such a proof would seem to be easily derivable from what is known about curves, corners, and inflection points surviving projection from three dimensions onto two.

Nonrelatable displays were obtained from relatable ones by shifting one piece in depth. In our experiments, depth shifts were accomplished by equal and opposite 2-D shifts in the two eyes. These monocular shifts induced small monocular misalignments in the image given to each eye: 5.6 arcmin of misalignment in the small shift condition and 11.2 arcmin of misalignment in the large shift condition. What are the effects on 2-D interpolation of these monocular shifts? Figure 10 shows the effects of these monocular shifts on the 4.5-cm depth-shifted and 9.0-cm depth-shifted displays. The displays in the figure have the same proportions as the experimental displays, and they subtend similar visual angles if viewed from a distance about 16 times the height of the white central figures in each display. It can be seen that the monocular shifts are small; phenomenologically, they do not seem to disrupt illusory contours, especially in the 4.5-cm depth shifted case.

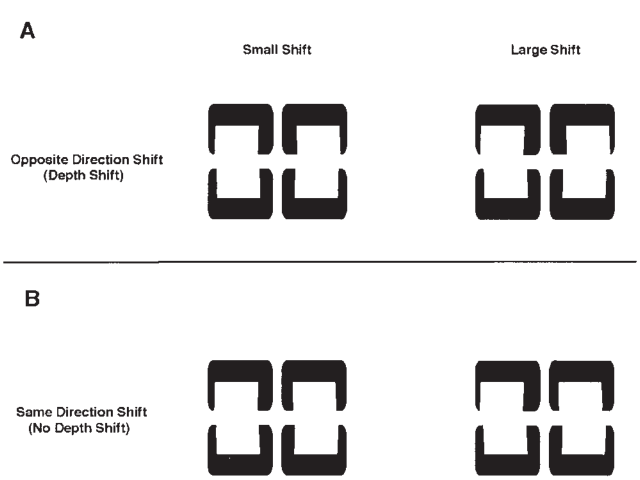

Figure 10. Same- and opposite-direction misalignment of left- and right-eye images. Each display is a stereo pair that can be free-fused by crossing the eyes. A: Shifted displays from Experiment 1, with opposite-direction shifts in the left and right eyes. B: Same-direction-shifted displays from Experiment 2; each eye's image had a magnitude of shift identical to the corresponding opposite-direction-shifted displays from Experiment 1.

What is known experimentally about the tolerance of contour interpolation for 2-D (planar) shifts? The geometric definition of relatability (Kellman & Shipley, 1991) actually indicates that any misalignment of parallel (or roughly parallel) edges falls outside the definition; however, Kellman and Shipley (1991) noted that any geometric condition on perceptual organization will be subject to thresholds. Shipley and Kellman (1992a) tested the effects of misalignment on completion; their data indicated that completion has a small but real tolerance to misalignment: Increasing amounts of misalignment of initially collinear edges reduces interpolation strength to nearly 0 by about 15 arcmin of visual angle. This value has held up surprisingly well in other studies. Kellman et al. (1998) used the fat–thin method developed by Ringach and Shapley (1996) and found that misalignment exceeding 15 arcmin also disrupted interpolation in their displays. Experiments by Punzel, Yonas, and Schrater (2001) using a different paradigm produced a similar estimate.

What is surprising about the consistency of this estimate is that, we believe, a particular amount of misalignment on the retina is unlikely to be the relevant determinant of object completion (Kellman et al., 2001). It is more likely that misalignment effects will depend on an angular notion based on a ratio—namely, the ratio of the horizontal misalignment of edges to their vertical separation. The convergence of estimates around 15 arcmin of visual angle may be an artifact of the use of displays with similar proportions across several experiments. This issue is currently under study. For the present purposes, the displays were designed explicitly so that the 2-D shifts used would fall within the tolerance expected for 2-D relatability, either on absolute retinal or figure-relative angular grounds. It was expected that the relatively small 2-D misalignments would not disrupt monocular relatability but would exert their effects as a result of the depth shifts created by the opposite misalignments in the two eyes. We cannot assume, however, that this was the correct assessment. It is possible that 2-D misalignments in each eye reduced the interpolation effects sufficiently to disrupt interpolation, producing the relevant differences between relatable and shifted displays in Experiment 1. To assess this possibility, we carried out Experiment 2.

Experiment 2

How can it be determined whether monocular misalignment, rather than the 3-D effects of misalignment in the two eyes, was responsible for the effects of relatable and nonrelatable displays? We used the following technique. In the 3-D displays, nonrelatable displays had equal and opposite misalignments in the two eyes. We made a new set of displays that had identical misalignments in the two eyes—that is, misalignments that were equal and in the same direction. An example of these displays, along with the earlier displays, is shown in Figure 10.

The displays with misalignments in the same direction induced no depth shift of one tab relative to the other. As can be seen by free-fusing the stereo pair in Figure 10B, these displays retained the slant of the tabs shown in Figure 5, but they had no depth shifts in the relation of the tabs, only planar shifts. Given that nonrelatable displays in Experiment 1 were based on depth-shifting from the corresponding relatable displays, the manipulation in this experiment aimed to preserve 3-D relatability while testing the effects of the same magnitudes of monocular (planar) misalignment used previously. The experimental question was whether the small 2-D shifts would disrupt 2-D relatability. On the one hand, if relatable displays maintained their advantage over nonrelatable displays in this experiment, it would suggest that the effects in Experiment 1 did not necessarily involve 3-D completion. On the other hand, if equivalent amounts of 2-D misalignment, without the 3-D shifts, did not produce the effects observed in Experiment 1, it would provide strong evidence that those effects were indeed 3-D in nature.

Method

All aspects of the method in Experiment 2 were the same as in Experiment 1, except as noted below.

Subjects. Subjects were 17 UCLA undergraduate students who received course credit for participation. All had normal or corrected-to-normal vision, and all passed a basic test for stereoscopic depth perception. An additional 2 subjects were not included in the final data set due to failure to meet a threshold criterion for performance on the experimental task (average d' [across conditions] $< .50$).

Stimuli. In this experiment, displays were either relatable or laterally shifted transforms of relatable displays. (We use the descriptor *laterally shifted* rather than *nonrelatable* to describe these because the experimental question was, in part, whether the misalignments used would disrupt 2-D relatability.) The misalignment stimuli were generated in the same manner as the stimuli from Experiment 1, with the exception that the disparity differences used to induce the depth shift were given in the same magnitude and direction in the two eyes rather than opposite directions. In other words, these misaligned stimuli had exactly the same amounts of misalignment as their corresponding 3-D nonrelatable stimuli in Experiment 1, but here the direction of misalignment was always identical in the two eyes. The amounts of lateral misalignment previously used to obtain the 4.5-cm and 9.0-cm depth-shifted displays were 5.6 arcmin and 11.2 arcmin, respectively. We refer to the two shift levels as the *smaller* and *larger* monocular shifts.

On the depth-completion hypothesis, shift conditions were predicted to produce results similar to those produced by the relatable displays in this experiment. The smaller monocular shift is clearly within the bounds of planar misalignment tolerated in earlier studies of 2-D interpolation, whereas the larger shift of 11.2 arcmin approaches but does not quite reach the 15-arcmin shift shown previously to disrupt 2-D relatability.

Results

The main results of Experiment 2 are shown in Figures 11 and 12. Figure 11 plots sensitivity as a function of stereoscopically specified tab slant for relatable, 4.5-cm shifted, and 9.0-cm shifted displays. In this experiment, sensitivity on the 3-D-relatable and shifted displays did not differ. Neither the smaller nor the larger monocular displacement reliably disrupted performance. Sensitivity as a function of slant resembled that observed for the relatable displays in Experiment 1 (see Figure 8), suggesting that interpolation occurred for all displays in Experiment 2.

These results were confirmed by the analyses. Sensitivity was analyzed in a 3 (lateral shift: unshifted, smaller shift, larger shift) × 4 (slant) within-subject ANOVA. There was a strong main effect of slant, $F(3, 14) = 53.93$, $p < .0001$. There was also a reliable main effect of shift, $F(2, 15) = 6.16$, $p < .05$. Contrary to

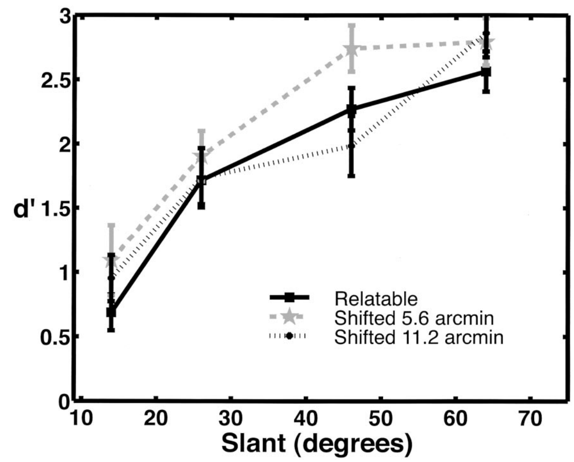

Figure 11. Sensitivity (d') as a function of slant in Experiment 2. Error bars represent plus or minus 1 standard error of the mean.

the results of Experiment 1, however, performance with the unshifted displays was not best. Averaged over slants, the d' means for unshifted, smaller shift, and larger shift displays were 1.81, 2.13, and 1.88, respectively. Individual comparisons showed that performance with relatable displays was actually worse than it was with the 5.6-arcmin shifted displays, $t(16) = -3.28$, $p < .01$, and did not differ from 11.2-arcmin shifted displays, $t(16) = -.63$, *ns*. The two levels of shift also differed, $t(16) = 2.33$, $p < .05$. Inspection of individual patterns showed that 13 of 17 subjects had higher overall d's for 5.6-arcmin shifted displays than for relatable ones, and 11 of 17 had higher overall d's for 11.2-arcmin shifted displays than for relatable ones.

There was also a reliable Shift × Slant interaction, $F(6, 11) = 3.45$, $p < .05$, due to the comparatively low performance on the unshifted displays at the smallest slant ($d' = 0.69$) and the comparatively high performance on the smaller shifted displays at the next to largest slant ($d' = 2.74$).

We also compared the pattern of results in this experiment directly with the pattern in Experiment 1. A 2 (experiment) × 3 (relatability) × 4 (slant) ANOVA was carried out, with experiment as a between-subjects factor and relatability and slant as within-subject factors. The analysis revealed a main effect of slant, $F(3, 30) = 85.12$, $p < .0001$. There was a main effect of experiment, $F(1, 32) = 7.06$, $p < .02$, due to the fact that in Experiment 2, all three shift conditions showed high levels of performance, whereas in Experiment 1, the depth-shifted conditions were substantially worse. This pattern was also reflected in the strong Experiment × Shift interaction, $F(2, 31) = 25.37$, $p < .001$. Individual comparisons showed that relatable (unshifted) displays in Experiments 1 and 2 did not differ in sensitivity, $t(32) = .54$, *ns*. However, sensitivity was higher in the laterally shifted displays of Experiment 2, relative to the depth-shifted displays of Experiment 1, for both 5.6-arcmin shifted, $t(32) = 3.77$, $p < .001$, and 11.2-arcmin shifted displays, $t(32) = 3.47$, $p < .002$. There was a reliable Slant × Shift interaction, $F(6, 27) = 3.39$, $p < .05$, probably reflecting the greater differences between experiments at higher slants in the shifted conditions. There was no reliable three-way interaction of experiment, shift, and slant. The advantage of later-

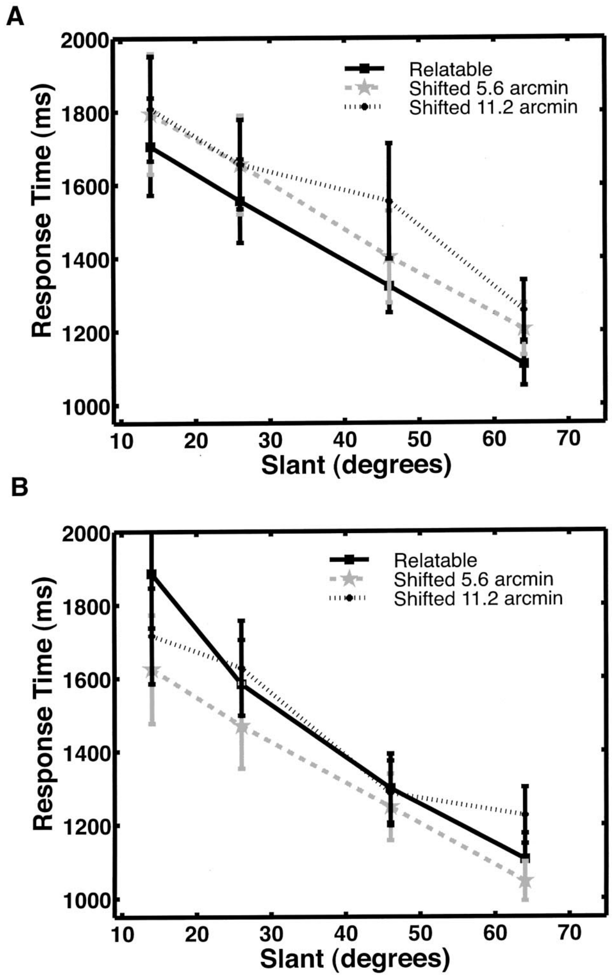

Figure 12. Response time as a function of slant for parallel (A) and converging (B) displays in Experiment 2. Error bars represent plus or minus 1 standard error of the mean.

ally shifted over comparable depth-shifted displays held for every value of slant and shift.

RT data were analyzed in a 3 (lateral shift) × 4 (slant) × 2 (display type: parallel, converging) within-subject ANOVA. There was a strong main effect of slant, $F(3, 14) = 15.04$, $p < .001$, and a small effect of display type, $F(1, 16) = 4.67$, $p < .05$—the latter produced by slightly faster overall responding to converging displays. There was also a reliable effect of shift, $F(2, 15) = 17.38$, $p < .001$. The only reliable interaction was Shift × Display Type, $F(2, 15) = 5.52$, $p < .05$.

These effects may be understood by examining Table 1. Because there were no reliable interactions with the slant variable, RTs are shown for shift condition and display type averaged over the several slants used. As suggested by Table 1, there was no reliable difference between the relatable and the smaller shifted displays, $t(16) = .385$, *ns*. The larger shifted displays produced longer RTs than both the relatable, $t(16) = 3.10$, $p < .01$, and the smaller shifted displays, $t(16) = 6.09$, $p < .001$.

In a further analysis, we compared the pattern of RTs in Experiment 2 directly with the pattern in Experiment 1. A 2 (experiment) × 3 (shift) × 4 (slant) × 2 (display type) ANOVA was carried out, with experiment as a between-subjects factor and shift, slant, and display type as within-subjects factors. There was a main effect of experiment, $F(1, 32) = 5.60$, $p < .03$; overall RTs were shorter for the subjects in Experiment 2 than for those in Experiment 1. There were reliable main effects of slant, $F(3, 30) = 27.83$, $p < .001$; shift, $F(2, 31) = 15.80$, $p < .001$; and display type, $F(1, 32) = 5.41$, $p < .03$. There was also a strong Experiment × Shift interaction, $F(2, 31) = 15.80$, $p < .001$, indicating, as with the sensitivity data, the effects of lateral shifting in Experiment 2 versus depth shifting in Experiment 1 for the nonrelatable displays. There were also several other small but reliable interactions: Display Type × Slant, $F(3, 30) = 4.15$, $p < .02$; Display Type × Relatability, $F(2, 31) = 3.51$, $p < .05$; Slant × Relatability, $F(6, 27) = 2.77$, $p < .05$; and Display Type × Slant × Experiment, $F(3, 30) = 3.75$, $p < .03$.

We followed up the effects of most importance for the primary experimental questions. Specifically, individual comparisons were used to examine the interaction between shift and experiment, which did not interact with display type or slant. The difference between relatable displays and each level of lateral shift (Experiment 2) was compared with the corresponding difference between relatable and depth-shift displays (Experiment 1). Results showed that the difference between relatable and the smaller shift displays in Experiment 2 (9.9 ms) was reliably smaller than the corresponding difference in Experiment 1 (219.7 ms), $t(32) = 4.41$, $p < .001$. This was also true for the comparison of relatable and the larger shift displays, in which the difference observed in Experiment 2 (95.7 ms) was reliably smaller than the difference in Experiment 1 (280.4 ms), $t(32) = 2.79$, $p < .01$. As with the sensitivity measure, the advantage of laterally shifted over comparable depth-shifted displays held for every value of slant and shift.

Discussion

The results of Experiment 2 clearly indicate that the 2-D misalignments, corresponding to those used in Experiment 1, are not sufficient to produce the differences between 3-D-relatable and nonrelatable displays observed in the earlier experiment. Experiment 1 showed large sensitivity and RT advantages of 3-D-relatable displays over nonrelatable ones; these advantages did not

Table 1
Response Times for Each Shift and Display Type in Experiment 2

Shift and display type	Response time (ms)
Relatable	
Parallel	1,423.3
Converging	1,419.1
Mean	1,421.2
Smaller monocular shift	
Parallel	1,514.8
Converging	1,347.5
Mean	1,431.5
Larger monocular shift	
Parallel	1,569.0
Converging	1,464.3
Mean	1,516.6

appear in Experiment 2. The displays in Experiment 2 were identical to those in Experiment 1, except that they were presented with same-direction shifts in each eye (instead of opposite-direction shifts to create depth differences). This manipulation not only eliminated the sensitivity advantage for relatable displays—surprisingly, it modestly reversed it. The 5.6 arcmin shifted condition actually showed slightly greater sensitivity than did unshifted displays in this experiment. Moreover, the levels of sensitivity in Experiment 2 for unshifted and shifted displays closely resembled sensitivity for the 3-D-relatable displays of Experiment 1. This suggests that all of these displays supported object completion.

For the smaller (5.6-arcmin) shift, the RT data likewise show no evidence of a superiority for unshifted displays. Unlike Experiment 1, RTs for unshifted and 5.6 arcmin shifted displays did not differ. The larger shift displays in Experiment 2 did show somewhat longer RTs (by about 100 ms). Interpretation of these data is complicated slightly by what appears to be a speed–accuracy trade-off. As the individual-subject data show, a majority of subjects were actually slightly less accurate on the unshifted displays than on the shifted ones for both levels of shift. For reasons that are unclear, subjects showed a bias to respond faster to unshifted displays in Experiment 2, at the cost of somewhat reduced accuracy. Despite this nuance, when accuracy and speed are evaluated together, it is clear that Experiment 2 did not reproduce the performance differences between 3-D-relatable and nonrelatable displays seen in Experiment 1. Moreover, sensitivity in the shifted conditions of Experiment 2 closely approximated those seen in Experiment 1 for 3-D-relatable displays.

Overall, the results of Experiment 2 provide evidence that the effects of nonrelatable displays observed in Experiment 1 depended on the depth shifts created by misaligning tabs in opposite directions in the two eyes. When the same magnitudes of shift were used but the direction was made the same in the two eyes, the difference between relatable and nonrelatable displays no longer produced higher sensitivity in our classification task. The simplest interpretation of this result is that the nonrelatable displays in Experiment 1 exerted effects on performance that were specifically results of their perceived positions in depth.

The results confirm those of earlier work (e.g., Kellman et al., 1998; Shipley & Kellman, 1992a) in indicating that misalignments of roughly parallel edges must exceed about 15 arcmin of visual angle to fully break up contour relatability in the plane. (As we indicated, this regularity may depend on stimuli having certain sizes and proportions; the actual constraint may be based on a ratio of misalignment to separation.) If a bound of approximately 15 arcmin of misalignment holds for these displays, we would expect that misalignments clearly exceeding that amount would break up object completion in the displays. For reasons described below, this prediction was tested in Experiment 3.

Experiment 3

Our results so far have been consistent with the idea that contour interpolation is a 3-D process and that it is governed by a smoothness constraint, formally given as 3-D relatability. This interpretation explains the strong superiority of sensitivity and speed for relatable displays in Experiment 1. It is also consistent with the results of Experiment 2, which showed that the effects could not be explained by 2-D misalignments.

Recall, however, our earlier discussion regarding alternative explanations. It is possible that the geometric relations given by relatability may just by coincidence be good ones for making comparisons of separate tabs. If this is so, then relatable displays would be classified more efficiently than nonrelatable ones, even without any 3-D object interpolation occurring.

A subset of the data in Experiment 1 was used to disconfirm the most salient member of this class of hypothesis, namely that comparisons involving tabs at similar depths might be easier than comparisons of tabs at differing depths. More subtle versions of this fortuitous-geometry hypothesis are not excluded by that analysis, however. The most subtle, general, and vexing version is that 3-D relatability does indeed make a difference but not because it produces a connection between tabs. The conditions that we have formulated as 3-D relatability might simply be, by coincidence, fortuitous spatial positions for performance of the comparisons of tabs required by the parallel–converging classification task.

In short, we are left with two possible explanations: Under the 3-D completion hypothesis, the effects are due to the visual system's connecting of visible parts into a unitary object. Under the fortuitous-geometry hypothesis, classification is enhanced for two separate parts simply by virtue of their relative positions in space. Our data suggest the importance of the geometry given by relatability in either case, but only on the former hypothesis does relatability describe the 3-D connecting process used to form objects.

Each of the next two experiments addressed this hypothesis while also addressing another substantive issue regarding 3-D interpolation. In an assessment of the fortuitous-geometry hypothesis, each constitutes a part of the same overall strategy, which rests on the following premise: Besides the geometry of contours and surfaces in 3-D space, the process of object completion has other requirements. If these requirements are not fulfilled, object completion is not predicted to occur, even though the visible areas meet the geometric requirements of 3-D relatability. If, however, the geometric relation of two pieces simply facilitates their comparison (rather than their connection into an object), then the presence of this geometry alone should govern the performance advantage observed with relatable displays.

What aspects of object completion are required, apart from certain 3-D spatial relations of their bounding contours? One is suggested by Experiment 2. Besides the proper orientational and positional relations in depth, contour-completion processes require that contours meet the criteria of 2-D relatability; constraints regarding relations in the observer's image plane are part of 3-D relatability, along with the specifically 3-D aspects. In Experiment 2, we found that the lateral misalignments used to create depth shifts were not large enough to disrupt interpolation—that is, they fell within the small tolerance of 2-D relatability for edge misalignment. We interpreted this result to mean that the effects depended on object completion and that completion was not disrupted by the amounts of lateral misalignment used. A different possibility, however, is that no amount of lateral misalignment would have disrupted the effects because the effects depended only on the depth relations of the tabs and not on any completion process. If this idea is correct, then using a greater amount of misalignment—one sufficient to disrupt 2-D relatability—should have little effect. If, however, object completion is critical, then more extreme amounts of lateral misalignment should destroy the advantage of the displays that are otherwise 3-D relatable over those that are not. These issues are addressed in Experiment 3.

Finally, an important regularity about contour interpolation is that interpolated contours in general begin and end at sharp corners (points of TD) in the optical projections of scenes. Kellman and Shipley (1991) presented a proof that instances of occlusion of one object by another would (except for degenerate cases) give rise to TDs in images. This invariant makes it difficult to study the role of TDs in occlusion displays, because they cannot be removed. It is possible, however, to remove TDs from illusory-contour displays. Shipley and Kellman (1990) showed that rounding of TDs in illusory-contour displays eliminated or drastically reduced the strength of contour interpolation (cf. Hoffman, 1998). In Experiment 4, we tested whether the geometry of 3-D relatability allows efficient classification when TDs have been smoothed to disrupt object completion.

The use of a large amount of misalignment in Experiment 3 also had a separate experimental purpose. As mentioned earlier, two complementary processes operate in object formation: a contour-interpolation process that creates units by making contour connections across gaps and a surface-spreading process that unifies surface features within real and interpolated contours (Grossberg & Mingolla, 1985; Yarbus, 1967; Yin et al., 1997, 2000). Surface spreading in 2-D illusory-contour displays has most often been discussed with regard to certain subsets of such displays, those that involve so-called "neon color spreading" or transparency (e.g., Bressan, Mingolla, Spillman, & Watanabe, 1997; Grossberg & Mingolla, 1985; Nakayama, Shimojo, & Ramachandran, 1990). One reason for this emphasis has been that spreading may be relatively weak in ordinary Kanizsa-style 2-D illusory-contour displays because all surfaces lie in the same depth plane. Thus, in the absence of contour interpolation to segregate surface regions, there is little surface spreading observed in displays such as that shown in Figure 13C. Likewise, the display in Figure 13A is consistent with surface spreading due to transparency, but because of depth ambiguity, occlusion may be seen instead. In Figure 13B, however, stereoscopic depth information has been used to separate depth planes, allowing a vivid impression of transparency. Similarly, surface spreading is more evident in Figure 13D than it is in Figure 13C, despite the absence of relatable contours in both cases.

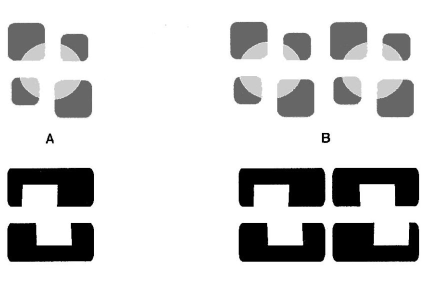

Figure 13. Relations between surface spreading and depth separation. A: Ambiguous transparency: The central object formed by interpolation may be perceived as transparent and in front of other surfaces or as partly occluded and behind other surfaces. B: Unambiguous transparency: The display is a stereo pair and can be free-fused by crossing the eyes. C: The display lacks relatable contours and shows minimal surface spreading. D: The display lacks relatable contours, but separation in depth from background shows obvious surface spreading. The display is a stereo pair and can be free-fused by crossing the eyes.

Because interpolated surfaces may occupy depth positions clearly different from the background and can curve in depth, 3-D interpolation may allow more obvious effects of surface spreading to be seen in illusory-contour displays such as that shown in Figure 13D. In this display, contours have been shifted to break up 2-D relatability, while surface spreading may still allow a smooth, monotonic connection between parts of the surfaces. This occurs because, even in the absence of closed forms given by real and interpolated contours, surface quality can spread beyond visible contour segments within their linear extensions (see Yin et al., 1997). There may appear to be some surface connection between the upper and lower tabs; however, there do not appear to be clear interpolated contours bounding the central region, and shape is, accordingly, somewhat amorphous. These observations suggest that surface interpolation in three dimensions may proceed without contour interpolation. This 3-D surface spreading may be similar to effects observed in structure-from-motion displays (Saidpour, Braunstein, & Hoffman, 1994).

The test of large misalignments in Experiment 3 was expected to shed light on the possibility of surface completion without contour interpolation. Assuming that large 2-D misalignments eliminated contour interpolation, as might be expected from prior research, we expected our results to reveal effects of the surface process alone on the parallel–converging discrimination task. On the one hand, if this task is sensitive primarily to effects of contour interpolation, we would expect little or no performance advantage for otherwise relatable displays with large 2-D shifts. On the other hand, if 3-D surface interpolation is sufficient for the performance advantages we have observed, we would expect such effects to be manifest with these displays as well.

Method

All aspects of the method in Experiment 3 were the same as in previous experiments, except as noted below.

Subjects. Subjects were 34 UCLA undergraduate students who received course credit for participation. All had normal or corrected-to-normal vision, and all passed a basic test for stereoscopic depth perception. An additional 5 subjects were not included in the final data set due to failure to meet a threshold criterion for performance on the experimental task (average d' [across conditions] $>$ .50).

Stimuli. Because of the lack of interactions of primary experimental effects with the slant variable in Experiments 1 and 2, we simplified the design to include only two slant values. The values chosen were the two most extreme used in previous experiments: the 14° and 64° slants. Stimuli in the control group were the same as used in Experiment 1 for these slant values. In the experimental group, these stimuli were all modified by the addition of a lateral misalignment of the top and bottom tabs of 30.7 arcmin.

Design. To ensure that the change from four to two slant values did not alter performance in our task in unsuspected ways, we tested both an experimental (large-misalignment) group and a separate control group in this experiment. The control condition was identical to Experiment 1, including both 3-D-relatable and nonrelatable displays, except for the use of only two slant values. Results from this control condition were intended to be compared with those from the experimental conditions of this experiment and of Experiment 4, in which two slant values were also used.

Results

The main results of Experiment 3 for sensitivity are shown in Figure 14. In the control condition, it is clear that, as in Experiment 1, sensitivity was higher for 3-D-relatable displays than it was for depth-shifted ones. In the misalignment condition, the difference between relatable and nonrelatable displays was reduced. Similar patterns appeared in the RT data (see Figure 15).

These observations were confirmed by the analyses. Sensitivity was analyzed in a 2 (condition: misalignment vs. control) × 3 (relatability: relatable, 4.5-cm shifted, 9.0-cm shifted) × 2 (slant) ANOVA, with condition as a between-subjects factor and relatability and slant as within-subject factors. There were main effects of relatability, $F(2, 64) = 24.4$, $p < .001$, and slant, $F(1, 32) = 247.16$, $p < .001$. There was no reliable main effect of condition, $F(1, 32) = 1.19$, *ns*.

The only reliable interaction was Condition × Relatability, $F(2, 64) = 3.56$, $p < .05$, reflecting the superiority of the control condition over the misaligned condition for 3-D-relatable displays but not for nonrelatable displays. This pattern was verified by two kinds of planned comparisons. We compared relatable displays with shifted displays within each condition. For the aligned (control) condition, sensitivity for 3-D-relatable displays was higher than that for 4.5-cm shifted displays, $t(16) = 4.27$, $p < .001$, and 9.0-cm shifted displays, $t(16) = 5.40$, $p < .001$. Sensitivity for the smaller shift was somewhat higher than for the larger shift, $t(16) = 1.85$, $p < .05$ (one-tailed). This pattern replicates the results of Experiment 1. The superiority of 3-D-relatable over shifted displays was observed for 16 of 17 subjects for the smaller shift and by all 17 subjects for the larger shift.

A different pattern was observed in the misaligned condition. Here, misaligned but otherwise relatable displays were not superior to 4.5-cm shifted displays, $t(16) = -.03$, *ns*. Both relatable, $t(16) > 7.72$, $p < .001$, and 4.5-cm shifted displays, $t(16) > 7.72$, $p < .001$, were superior to 9.0-cm shifted displays.

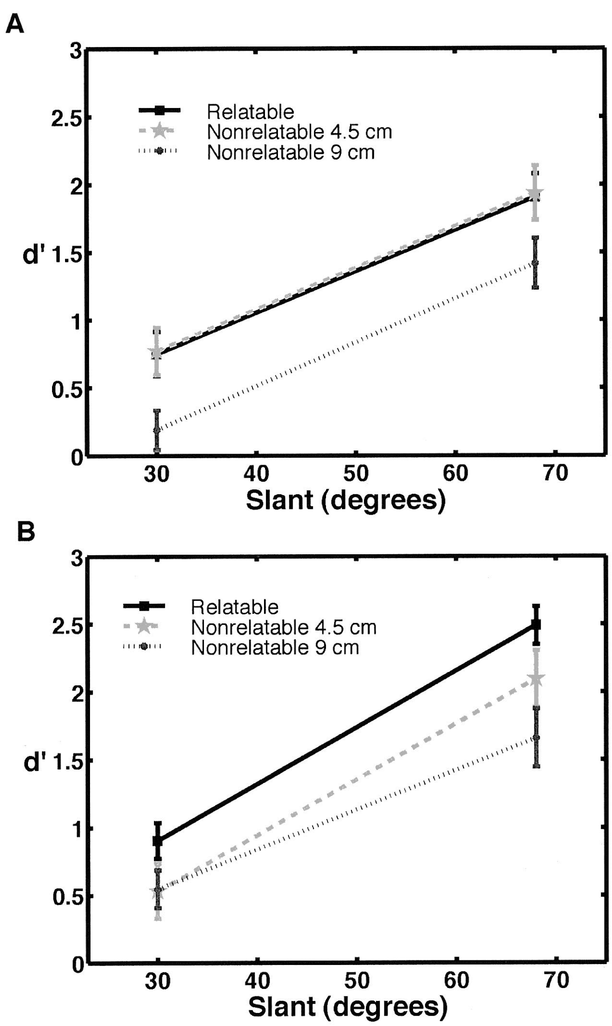

Figure 14. Sensitivity (d') as a function of slant for the experimental (planar misalignment) condition (A) and the control (planar alignment) condition (B) in Experiment 3. Error bars represent plus or minus 1 standard error of the mean.

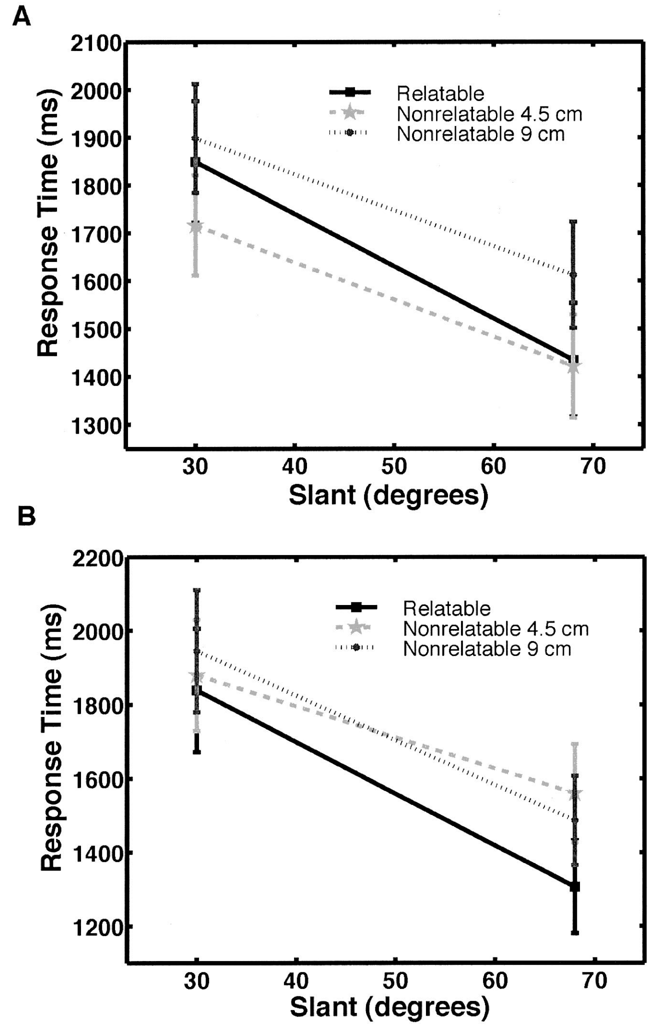

Figure 15. Response time as a function of slant for the experimental (planar misalignment) condition (A) and the control (planar alignment) condition (B) in Experiment 3. Error bars represent plus or minus 1 standard error of the mean.

We also compared each level of relatability across conditions. These comparisons showed that aligned 3-D-relatable displays produced reliably higher sensitivity levels than did misaligned relatable displays, $t(32) = 1.94$, $p < .05$ (one-tailed), but there were no reliable differences for displays having either a 4.5-cm depth shift, $t(32) = .08$, *ns*, or a 9.0-cm depth shift, $t(32) = 1.41$, *ns*.

RT as a function of slant is shown for the two conditions in Figure 15. In the control condition, 3-D-relatable displays were processed fastest, especially at the larger slant value (cf. Figure 9). In the misaligned condition, displays fitting the geometry of relatability (but laterally misaligned by 30.7 arcmin) did not show faster responses than depth-shifted displays.

These patterns were confirmed by the statistical analyses. RT was analyzed in a 2 (condition) × 3 (relatability) × 4 (slant) × 2 (display type: parallel vs. converging) ANOVA, with condition as a between-subjects factor and relatability, slant, and display type as within-subject factors. There were large main effects of slant, $F(1, 32) = 41.00$, $p < .0001$, and relatability, $F(2, 31) = 7.15$, $p < .01$. There was no reliable main effect of condition, indicating that misaligning the displays in one condition did not have much effect on the overall processing time required by the task. There was a reliable Relatability × Condition interaction, $F(2, 31) = 11.69$, $p < .001$, resulting from the superiority of relatable displays in the aligned condition to those in the misaligned condition. There were also Display Type × Condition, $F(1, 32) = 5.98$, $p < .05$; Display Type × Slant, $F(1, 32) = 9.92$, $p < .01$; and Slant × Condition interactions, $F(2, 31) = 5.622$, $p < .01$, indicating some variation across conditions in speed of processing for converging and parallel displays. Individual comparisons showed that in the aligned condition, 3-D-relatable displays were processed faster than both 4.5-cm shifted, $t(16) = 3.43$, $p < .01$, and 9.0-cm shifted displays, $t(16) = 3.41$, $p < .01$. There was no reliable difference between the two levels of shifted displays, $t(16) = 1.42$, *ns*. A different pattern appeared in the misaligned condition. Relatable displays produced marginally longer RTs than did 4.5-cm shifted displays, although the difference was not reliable, $t(16) = 1.50$, *ns*. The 9.0-cm shifted displays were marginally worse than the relatable displays, $t(16) = 1.79$, $p < .05$ (one-tailed), and reliably worse than the 4.5-cm shifted displays, $t(16) = 4.44$, $p < .01$.

Differences between 3-D-relatable and nonrelatable displays were compared between the aligned and misaligned conditions. These tests indicated that the difference between relatable and 4.5-cm-shifted displays was greater in the aligned than in the misaligned condition, $t(32) = 3.44$, $p < .002$. The difference between relatable and 9.0-cm shifted displays, compared across conditions, did not reach significance, $t(32) = .41$, *ns*.

Discussion

One goal of Experiment 3 was to compare the 3-D-completion hypothesis with the fortuitous-geometry hypothesis. The displays used in the misalignment condition were designed to disrupt object completion, using a factor separate from the depth relationships of the visible parts. If the earlier results depended on 3-D completion, we expected those effects to be disrupted by this lateral-misalignment manipulation. If the earlier results depended on some fortuitous benefit of depth relationships for performing slant comparisons, then the lateral misalignment should not have differentially disrupted displays that would otherwise produce object completion.

The results indicated that large lateral misalignments did disrupt performance. The pattern is consistent with the 3-D-completion hypothesis. In particular, the data showed an interaction such that lateral misalignment differentially affected 3-D-relatable displays. This pattern would be expected if lateral misalignment and depth shifting both disrupted object completion, because either one alone should have lowered performance. The combination, however, should not have lowered performance further. Lateral misalignment did not show a generally disruptive effect across all conditions. There was no reliable main effect of experimental condition in either the sensitivity or RT data. This outcome suggests that lateral misalignment is not an independent factor making classification harder in this experimental paradigm. Rather, it exerts its effects as a factor that can disrupt object completion that would otherwise occur in some displays.

Results in the control condition in this experiment replicated those in Experiments 1 and 2. Consistent with the absence of interactions with the slant variable in the earlier results, the use of only two slant values did not seem to make a difference in the results.

The second major goal of Experiment 3 was to assess the effect of 3-D surface interpolation on our experimental task. The results suggest that performance differences in the parallel–converging classification task that depend on interpolation rely more on contour processes than on surface processes. Clear superiority of 3-D-relatable over depth-shifted displays was seen in the control condition, but these effects were largely eliminated by large 2-D misalignments in the experimental condition, despite the fact that the misaligned displays probably induced some 3-D surface interpolation between the visible tabs.

To avoid possible confusion, it may be useful to comment on the relationship of Experiments 2 and 3. Both used lateral misalignment, yet their results differed conspicuously. Experiment 2 showed that the amounts of misalignment present in the monocular components of the original depth-shifted displays were not sufficient to disrupt object completion. Relative to Experiment 1, Experiment 2 introduced no new misalignments; rather, the opposite-direction misalignments used to create depth shifts were replaced by misalignments of the same magnitude and direction. These misalignments were well within the range of known tolerance of object completion to planar misalignment (in the case of the smaller, 5.6-arcmin misalignment) or near the borderline of that range (in the case of the larger, 11.2-arcmin misalignment). The results suggest that under these conditions, completion occurred in both relatable and monocular shifted groups.

Experiment 3 used a greater amount of lateral misalignment, one that should clearly have disrupted 2-D relatability and, therefore, reduced or eliminated object completion based on contour interpolation in the displays. This manipulation did disrupt completion, eliminating the performance advantage of otherwise 3-D-relatable over nonrelatable displays. Because large lateral misalignments are known to disrupt object completion, these results suggest that the advantages in the control condition of this experiment and in Experiment 1 depended on perceptual unification of visible pieces in 3-D-relatable displays. They implicate contour rather than surface interpolation as the basis of performance differences, and they favor the 3-D-completion explanation over the fortuitous-geometry explanation for the accumulated results of Experiments 1–3.

Experiment 4

Experiment 4 explored an important determinant of contour interpolation and, specifically, its role in 3-D object formation. Evidence suggests that interpolation normally begins and ends at points of TD in an image (Kellman & Loukides, 1987; Kellman & Shipley, 1991; see also Heitger, Rosenthaler, von de Heydt, Peterhans, & Kubler, 1992; Heitger et al., 1998). TDs are points along contours at which there is no unique slope. They include contour junctions of all types, because the sharp intersection of contours will ordinarily produce two orientations at their point of intersection. The regularity that contour interpolation begins and ends at TDs is connected to an ecological fact. In cases of partial occlusion of an object, TDs will generically be produced at the points of disappearance of contours of one object behind an occluding object (for a proof, see Kellman & Shipley, 1991, Appendix A). This fact makes it impossible in practice to test the effects of the presence or absence of TDs in ordinary partial-occlusion displays (because TDs are always present). It is possible, however, to manipulate the presence or absence of TDs in illusory-contour displays. Shipley and Kellman (1990) reported evidence that rounding of TDs eliminated illusory contours in some cases and greatly reduced them in others. Other research has supported this claim, although some researchers have argued that eliminating TDs weakens but does not eliminate interpolation, because second-order discontinuities in first-order continuous contours may also contribute to initiating interpolation, at least in some contexts (Albert, 2001). (A second-order discontinuity occurs, e.g., where a straight segment meets a constant-curvature segment and their slopes match at the join point.) The difference between cases in which illusory contours are eliminated and cases in which they are weakened may have to do either with the presence of second-order discontinuities or with other issues, such as the presence of junctions at different spatial scales. Guttman and Kellman (2002) suggested that residual evidence of interpolation in displays with rounded TDs may reflect the activity of junction detectors at low spatial frequencies (cf. Wurtz & Lourens, 1999).

Experiment 4 was designed to test 3-D contour interpolation in the absence of TDs. The displays in this experiment mirrored displays in Experiment 1 in terms of the relative positions and orientations of physically specified tabs. This manipulation served two primary experimental purposes. First, it represented the final piece of our strategy for testing the fortuitous-geometry hypothesis. On any hypothesis arguing that the earlier observed effects of 3-D relatability derived only from relatability's prescribing favorable positional relations for comparison of separate tabs (i.e., any version of the fortuitous-geometry hypothesis), relatability effects would have been expected to appear in this experiment. However, if 3-D relatability effects derive from their role in object formation, then the elimination of TDs should have removed the beneficial effects of relatability. This follows because, if TDs are necessary conditions for interpolation, their removal should weaken or eliminate interpolation between the visible parts.

Second, the experiment served to test the importance of TDs in 3-D interpolation. Although our general view of contour interpolation as a 3-D process suggests that TDs should be just as crucial as initiating conditions for interpolation in three dimensions as they are in two dimensions, we know of no prior experimental test of this conjecture. To obtain a clear-cut absence of TDs in our displays, we developed visible areas with highly rounded corners (see Figure 16); for convenience, we refer to these as *cloverleaf* displays.

Method

All aspects of the method were the same as in previous experiments, except as noted below.

Subjects. Subjects were 14 UCLA undergraduate students who received course credit for participation. All had normal or corrected-to-normal vision, and all passed a basic test for stereoscopic depth perception. An additional 7 subjects were not included in the final data set due to failure to meet a threshold criterion for performance on the experimental task (average d' [across conditions] $> .50$).

Stimuli. The cloverleaf stimuli had the same overall size and depth relationships as the tabs used in Experiment 1. The corners of these stimuli were rounded to eliminate TDs. Extreme rounding was done because stimuli with only slightly rounded corners may activate junction detectors at coarse levels of scale (Shipley & Kellman, 1990; Wurtz & Lourens, 1999). To enhance the detectability of the curvature of the stimuli, we increased contrast in this experiment by changing the stimulus colors. Cloverleaf stimuli appeared as dark green against a white background, producing a Michelson contrast of 89%.

Results

The main results of Experiment 4 are shown in Figures 17 and 18.

Sensitivity data. Figure 17 plots sensitivity as a function of stereoscopically specified slant for the unshifted and shifted conditions of the experiment. For both slants, sensitivity did not differ as a function of whether the displays fit the geometry of 3-D relatability, were shifted 4.5 cm, or were shifted 9.0 cm. These observations were confirmed by the analyses. Sensitivity was analyzed by a 3 (relatability: unshifted, 4.5-cm shifted, 9.0-cm shifted) $\times$ 2 (slant) ANOVA, with relatability and slant as within-subject factors. There was no reliable main effect of relatability, $F(2, 26) = 1.10$, *ns*, but a large main effect of slant, $F(1, 13) =$

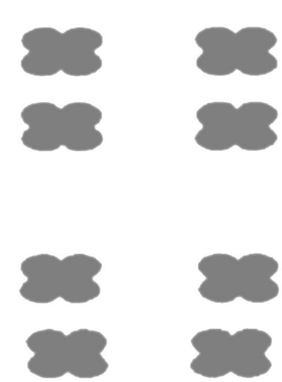

Figure 16. Examples of displays used in Experiment 4: A relatable (converging) display with rounded corners (top) and a nonrelatable (converging) display with rounded corners (bottom). Both displays are stereo pairs and can be free-fused by crossing the eyes.

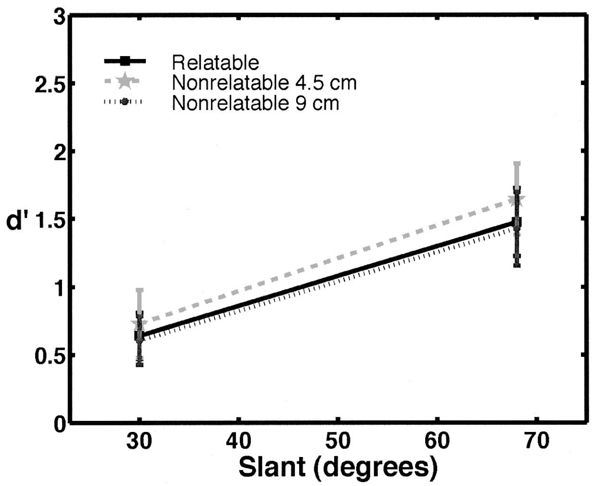

Figure 17. Sensitivity (d') as a function of slant in Experiment 4. Error bars represent plus or minus 1 standard error of the mean.

57.81, $p < .0001$. There was no reliable Relatability × Slant interaction, ($F < 1$).

Sensitivity data in this experiment were compared to identically positioned displays having TDs. Because two values of slant were tested, we compared the data with those of the control group in Experiment 3, which had the same two values of slant. This analysis of sensitivity comprised a 2 (condition: cloverleaf displays vs. regular displays) × 3 (relatability: unshifted, 4.5-cm shifted, 9.0-cm shifted) × 2 (slant) ANOVA, with condition as a between-subjects factor and relatability and slant as within-subject factors. There was no main effect of condition, $F(1, 29) = 1.37$, *ns*, indicating that cloverleaf displays were not harder or easier to classify in general than the display form used in earlier experiments. There were reliable main effects of relatability, $F(2, 28) = 7.47$, $p < .01$, and slant, $F(1, 29) = 201.11$, $p < .0001$. There were

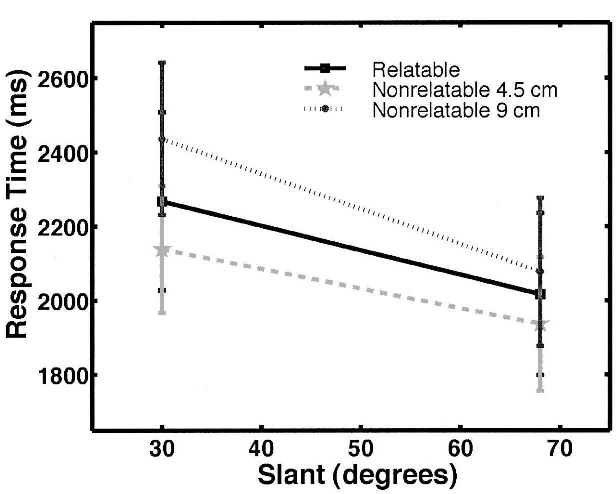

Figure 18. Response time (RT) as a function of slant in Experiment 4. RTs for parallel and converging stimulus types did not differ reliably and were combined. Error bars represent plus or minus 1 standard error of the mean.

also reliable Slant × Condition, $F(1, 29) = 12.36$, $p < .01$ and Relatability × Condition interactions, $F(2, 28) = 8.90$, $p < .001$.

The latter effect derived from the fact that the difference between 3-D-relatable and nonrelatable displays appeared only in the displays containing TDs. Averaged across the two slants, the overall d' values for the illusory-contour versus the cloverleaf displays were 1.70 versus 1.06 for unshifted displays, 1.31 versus 1.18 for 4.5-cm shifted displays, and 1.10 versus 1.02 for 9.0-cm shifted displays. This pattern indicates that only the relatable displays in the original stimulus set showed an advantage over nonrelatable displays. Within the cloverleaf condition, there were no reliable differences between unshifted displays and either level of shift (all ts[13] < 1.35, *ns*). Direct comparisons at each shift level showed that 3-D-relatable displays produced higher sensitivity in the displays with TDs, $t(29) = 2.83$, $p < .01$, whereas there were no differences between the groups for either 4.5-cm shifted, $t(29) = .43$, *ns*, or for 9.0-cm shifted displays, $t(29) = .35$, *ns*.

RT data. Similar patterns were shown in the RT data. Figure 18 shows RT as a function of slant for the several configurations. These data were analyzed in a 3 (relatability) × 2 (slant) × 2 (display type: parallel, converging) within-subjects ANOVA. There was a reliable effect of slant, $F(1, 13) = 8.76$, $p < .02$, but no other reliable main effects or interactions.

The pattern of RT results in this experiment was also compared directly with the pattern from the control group in Experiment 3 in a 2 (experiment) × 3 (relatability) × 2 (slant) × 2 (display type) ANOVA, with experiment as a between-subjects factor and relatability, slant, and display type as within-subject factors. The analysis showed a reliable main effect of experiment, $F(1, 29) = 4.53$, $p < .05$, due to the somewhat shorter RTs overall in the displays having TDs. There were also main effects of slant, $F(1, 29) = 26.5$, $p < .001$, and relatability, $F(2, 28) = 5.22$, $p < .02$, and a reliable Experiment × Relatability interaction, $F(2, 28) = 4.95$, $p < .02$. There were no other reliable main effects or interactions. This pattern indicates that whereas 3-D-relatable displays differed from shifted ones when displays contained TDs, they did not do so in the cloverleaf displays that had the same positions and orientations. We examined the interaction by finding within each group the differences between RTs for 3-D-relatable and shifted displays and comparing these differences across groups. (Because of the absence of any reliable interaction of these effects with slant, these data combined both slant values in each group.) Differences are shown by group and shift level in Table 2. The analysis showed that the difference between relatable and 4.5-cm shifted displays was reliably greater in the control group, $t(29) = 2.50$, $p < .02$. The difference between relatable and 9.0-cm shifted displays did not reach significance, $t(29) = .32$, *ns*.

Table 2
Comparison of Mean Response Times (in Milliseconds) by Display Type and Shift Magnitude in Experiment 4

Display type	Unshifted vs. 4.5-cm shift	Unshifted vs. 9.0-cm shift
Control	−147	−143
Cloverleaf	+105	−114

Discussion

The data from Experiment 4 clearly indicate that rounding of TDs eliminates the difference between displays whose visible parts fit the geometry of 3-D relatability and those whose visible parts do not. No advantage for relatable over nonrelatable displays was found, either in sensitivity or RT. Moreover, sensitivity for all displays fell well short of the levels of performance found in Experiments 1 and 2 for 3-D-relatable displays.

The information that these data provide about 3-D object formation is important for two reasons. These data unequivocally support the depth-completion hypothesis. Placing the visible tabs in particular geometric relations that satisfy the geometry of 3-D relatability did not facilitate performance when other requirements for object completion were lacking. The cloverleaf displays in this experiment shared identical 3-D positions and orientations with the relatable tabs in earlier experiments. Yet these orientations and positions were insufficient to produce the effects on sensitivity and speed observed in earlier experiments. Removing some other requirement of object completion, such as 2-D alignment or TDs, removes the effect of 3-D relatability.

The results of Experiment 4 confirm the importance of TDs as features that initiate object-completion processes in three dimensions. The present results agree with those of other research (Guttman & Kellman, 2002; Palmer, Kellman, & Shipley, 2000; Shipley & Kellman, 1990) that rounding of TDs reduces or eliminates interpolation. These data provide the first evidence we know of regarding the importance of TDs in 3-D interpolation.

Experiment 5

Experiments 1–4 validated an objective method for the study of 3-D interpolation, showed that 3-D interpolation occurs, and provided support for a geometric account—3-D relatability—of the conditions that lead to 3-D interpolation. Below (see the General Discussion), we elaborate on the implications of the results for a general account of object formation. To do that, we first address one other issue, important for the theory of 3-D contour interpolation and for relating our results to those from studies of 2-D interpolation and of slant perception.

Specifically, Experiment 5 addressed how much slant was perceived in the displays used in our experiments. The stimuli used in Experiments 1–4 used two or four slants whose values have been labeled by the amount of slant specified stereoscopically. Slant variation was used primarily as a carrier variable to examine task performance at varying levels of difficulty. The amounts of slant actually perceived by our subjects was not addressed in these experiments. Do the stereoscopically specified slants in our displays lead to similar perceived slants? The findings regarding 3-D relatability and its disruption by depth shifts did not depend on any particular slant-perception performance by subjects, except perhaps that more slant should have made the task easier, which was true in the results of all experiments.

In using the results of these studies to build a geometric account of 3-D interpolation, however, there is at least one issue for which the magnitude of perceived slant is important. Theory and research on 2-D object completion indicate that strength of interpolation between contours in two dimensions decreases as these contours depart from collinearity, becoming weak (Guttman et al., 2003) or absent (Field et al., 1993; Kellman & Shipley, 1991) when the orientation of two contours differs by 90° or more. In other words, interpolation is weak or absent for contours forming acute angles.

Does a similar principle of relative angle govern 3-D interpolation? In the present experiments, slants of the visible tabs ranged from 14° to 64° away from frontoparallel. In the converging displays, two tabs slanting away from frontoparallel at 64° would have an orientation difference from each other of 52°. For the four slants used in these experiments, the converging tabs formed angles of 152°, 128°, 88°, and 52°. If a 90° principle governs 3-D interpolation, then contours whose linear extensions form acute angles in 3-D space should not support interpolation. If this is so, should our two or three highest slant values have led to difficulties in the converging stimuli?

If slant was perceived veridically, the prediction of a 3-D 90° constraint would be that at the greatest slant magnitudes presented in this study, a depth shift would be redundant, because relatability would already be disrupted due to the angular relation of the illusory tabs. This would, of course, only be true in the converging displays; parallel displays would be unaffected. On this assumption, we would have expected decreased performance advantage for relatable converging displays over nonrelatable ones at the higher slant values. Such a pattern would have produced a Relatability × Slant interaction for sensitivity and a Relatability × Slant × Display Type interaction for RT in Experiment 1. No such reliable interaction occurred for sensitivity. There was a reliable three-way interaction for RT, but it went in the other direction: As can be seen in Figure 9, the advantage for converging, 3-D-relatable displays actually increased at the higher slant values.

These results are consistent with two possibilities. One is that a 90° constraint governs relatability in the plane but not in depth. The other possibility is that slant was consistently underperceived in our displays because of the lack of correlation between stereoscopic and pictorial information. If slant was substantially underperceived, the present results simply do not bear on the 90° constraint.

From the literature on slant perception, we would expect that slant given solely by stereoscopic disparity might be substantially underperceived (Gillam & Ryan, 1992; Howard & Rogers, 1995; van Ee, Banks, & Backus, 1999). Experiment 5 was undertaken to investigate slant perception in our experimental situation so as to determine what implications, if any, the current results have for a 90° constraint in 3-D interpolation.

Observers' casual reports of slant magnitudes, as well as our own observations, indicated that the perceived angular slants were significantly less than the slants specified in the stimuli. Experiment 5 attempted to confirm and quantify the underperception of slant that was believed to have occurred.

Method

Subjects. Two naive and two nonnaive subjects participated in this experiment. Two of the subjects were male, and two were female. All had normal or corrected-to-normal vision and were given a stereoacuity pretest. No subjects were excluded from the experiment on the basis of the results of the stereoacuity pretest. No subject received any incentive for participating in this study.

Apparatus. The stimuli in this experiment were presented in the same manner as Experiment 1. To measure perceived slant, we constructed an adjustment apparatus described by Gillam and colleagues (Gillam &

Blackburn, 1998; Gillam & Ryan, 1992) that has been used in numerous studies of slant perception. A thin circular wheel (from a Meccano [Clichy, France] building set) was rigidly connected to a thin metal rod inserted into a rectangular black housing (see Figure 19). This allowed the Meccano wheel to pivot about its horizontal diameter. The wheel, rod, and housing resided inside an illuminated box, painted white on the inside and black on the outside. The design of the box was intended to allow the disk to be lit while minimizing ambient light in the room. This setup was intended to match the conditions of the previous four experiments in which the room was kept dark. The disk was positioned at the subject's eye level, which was equal to the height of the stimuli that were presented on the screen, and to the left of the screen relative to the observer. The adjustment apparatus was mounted on a horizontal rotating dais, allowing subjects to vary the wheel's orientation somewhat to obtain additional perspective information. Subjects manipulated the disk manually with their left hands, under binocular viewing, using the wheel itself or a thin metal arm that connected perpendicularly to the rotating horizontal rod (see Figure 19). Because of the symmetric pattern of apertures in its surface, the Meccano wheel provides rich monocular perspective cues that allow accurate slant perception (Gillam & Ryan, 1992).

True slant of the wheel was recorded after each trial. The slant of the wheel was automatically converted to degrees through the output of a potentiometer within the black housing, attached to the pivoting wheel. Before each subject began the experiment, the slant estimator was calibrated at 5° increments through all possible slant positions of the wheel. Each position was sampled twice, and the average of the two estimates was used as the true slant at that position. A linear interpolation between each pair of measured slants was used to estimate all intermediate slant values. A comparison of the calibrations from each subject indicated that the slant estimator's properties did not vary throughout the experiment.

Stimuli. The displays from Experiment 1 were used. A subset of these displays was modified by removing from each display either the top or bottom tab and background element. In addition to the complete set of stimuli from Experiment 1, every top and bottom tab and inducing element that was presented as part of a pair in Experiment 1 was presented individually in this experiment.

Procedure. Subjects in this experiment were presented with each of the stimuli and instructed to match the slant of the illuminated disk to the slant of the observed tab using the slant-matching apparatus. In the case in which two tabs were presented, subjects were instructed to match the tab to the slant of the top tab. When satisfied that the slants of the relevant surface and the disk were matched, the subject pressed a red button on top of the black housing to record the value of the wheel's slant, and the next stimulus was presented.

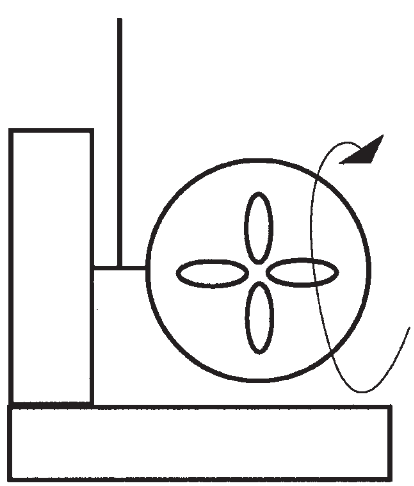

Figure 19. Schematic of the adjustment wheel used to measure perceived slant in Experiment 5.

Results

Results indicated that subjects systematically underperceived the slant magnitudes of the tabs relative to the values specified through stereodisparity. Figure 20 depicts the main results. The data shown are typical of the results for all 4 subjects. Despite considerable variability in the subjects' estimations of angular slant using the Meccano wheel, a linear fit to each subjects' data indicated that the underperception was reliable ($p < .05$). However, significant individual differences were also observed. Underperception ranged from 38% to 60%. Mean estimates for the four slants used in the experiment (which had virtual slant values of 14°, 26°, 46°, and 64°) were 11°, 15°, 21°, and 29°, respectively (averaged over positive and negative orientations for each slant). Maximum values of any subject's mean were 14°, 17°, 25°, and 35° for the four slant conditions. Subjects were reasonably consistent with each other; standard deviations for the subject means for the four slants were 7.1°, 6.3°, 5.3°, and 5.9°, respectively. Although we did not have enough subjects for a statistical test, underperception did not appear to covary with stereoacuity. No effects of tab position (top vs. bottom in single tab presentation trials) or of number of tabs presented (single vs. double tab trials) were observed (all ps $> .05$).

Discussion

The results from Experiment 5 indicate that slant was underperceived in the displays used in Experiments 1–4. Such underperception is common in slant-perception experiments (Gillam & Ryan, 1992, provided an extensive study of potential causal factors in the underperception of slant; see also Howard & Rogers, 1995). One consistent finding is that slant will be underperceived when perspective cues do not correlate with slant. In our experiments, the projective widths and heights of the tabs were always held constant. This served several purposes of experimental control, but it also precluded correlation of the perspective information with stereoscopic depth information. Moreover, although one might interpret the absence of any perspective variation to mean that the displays contained no perspective cues, it is also possible that keeping display proportions constant acted as a competing cue, providing some information that the displays were frontoparallel in every case. The effect may have been especially strong for displays near frontoparallel. Our data suggest that variability did not scale with mean slant value as might have been expected; rather, the largest standard deviations were shown for the smallest slant value. In sum, the data from this experiment indicate that the previous experiments did not provide a test of angular limitations (e.g., the 90° bending limit) for contour interpolation in three dimensions. As shown in Figure 20 and the subjects' means, all of the slants used in the earlier experiments led to perceived 3-D slants that fell well within a 90° limit for most observers. (Perceived slant of 45° for each single tab would produce a 90° interpolation path for a converging display.) The largest subject mean observed for the maximum slant condition (or any slant condition) in Experiment 5 was about 35°.

In 2-D interpolation, studies have indicated that the strength of interpolation decreases as the interpolated contour becomes more

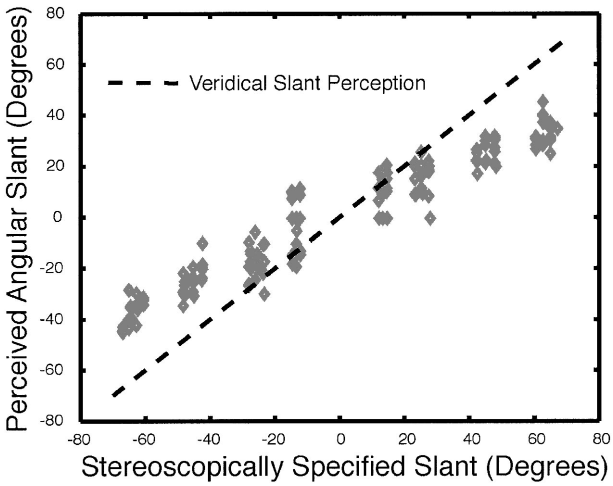

Figure 20. A representative set of data points from Experiment 5, plotting perceived slant (measured by adjustment) as a function of stereoscopically specified slant.

curved and that configurations in which an interpolated contour must bend through 90° or more have very weak or no interpolation (Kellman & Shipley, 1992). Some recent results, however, are consistent with the possibility that 2-D interpolation can exceed a 90° limit in some circumstances, although interpolation requires more time (Guttman et al., 2003). The current results do not provide much information about an angular limit in 3-D interpolation. The results of Experiment 5 suggest that the underperception of slant in these experiments was large enough for most subjects that the perceived slant magnitudes of almost all of the converging stimuli were effectively well within the boundaries where interpolation should be robust.

General Discussion

Taken together, these experiments support several important conclusions about object perception. Most generally, contour interpolation processes are 3-D in nature. Not only do they produce unified representations from retinally separated visible areas, they also take as inputs positions and orientations of edges in 3-D space. The outputs of these processes are interpolated contours and surfaces that extend through all three spatial dimensions.

These conclusions follow from the experiments. Experiment 1 tested our notion of 3-D relatability by comparing 3-D-relatable displays with displays in which one visible part was shifted in depth relative to the other to disrupt relatability. The effects of these positional relations on object formation were studied using a classification task in which subjects judged the relative orientations of two planar tabs. Although correct classifications on the task did not depend on whether displays were relatable or nonrelatable, we hypothesized that object formation would produce advantages in classification. Large advantages in both sensitivity and RT were found for relatable displays relative to both the 5.6 arcmin and 11.2 arcmin disparity shifts used to disrupt relatability.

Experiment 2 showed that the effects in Experiment 1 could not be explained by monocular shifts of contours used to create stereoscopic depth shifts. Classification performance for displays that had the same small monocular shifts as those in Experiment 1 but did not have the depth shifts was far better than it was for shifted displays in Experiment 1, and it was similar to performance for the 3-D-relatable displays in Experiments 1 and 2.

Analysis of a subset of stimuli in Experiment 1 showed that a simple alternative hypothesis—that having tabs at more nearly equal observer-relative depths—could not explain the advantage of 3-D-relatable displays. However, this alternative explanation was just one of a larger class of hypotheses that needed to be considered. The geometric relations given by 3-D relatability might have facilitated classification performance either because they led to unitary object formation (the object-completion hypothesis) or because they just happened to be helpful relations for making comparisons of separate objects (the fortuitous-geometry hypothesis).

Experiments 3 and 4 used converging methods to rule out the fortuitous geometry hypothesis and produce new information about 3-D interpolation. Each experiment used a different manipulation designed to reduce or eliminate object formation while keeping the depth relations of visible tabs the same as in Experiment 1. The large 2-D misalignments in Experiment 3 largely removed the advantages of relatable over nonrelatable displays, as would be expected if both lateral and depth aspects of 3-D relatability are required for object formation. Moreover, because the lateral-shift manipulation in Experiment 3 disrupted contour interpolation but did not remove the conditions for surface spreading, the results indicated that classification advantages for relatable displays in earlier experiments derived primarily from processes of contour interpolation, not surface interpolation. Experiment 4 used displays that preserved the 3-D positions for relatable and shifted displays, but displays in all cases had rounded corners to eliminate TDs. If TDs are crucial features for initiating 3-D interpolation, and if the superiority of 3-D-relatable over nonrelatable displays derives from object interpolation, then it was expected that with rounded displays, performance on relatable displays would fall to the level of that on nonrelatable ones. This was exactly what happened. The effects of relatability on classification derive from object formation; disrupting other crucial ingredients for object formation, such as the presence of TDs, removes the effect of relatability.

Finally, Experiment 5 asked whether the present data can be used to address the operation of a 90° constraint in 3-D interpolation. We quantified the effects of slant underestimation in the experimental displays and found that it was reasonably consistent across subjects. The general underestimation of slant has little consequence for the findings about the effects of 3-D relatability and disruptions of relatability due to depth shifts. However, the results do indicate that it is likely that none of the experimental displays violated the 90° principle in three dimensions, and thus, the superior classification performance for relatable displays at all slants tested does not disconfirm the possibility that a 90° principle applies in three dimensions. These results go beyond earlier observations about 3-D illusory-contour formation to place 3-D interpolation on a firmer experimental footing. Specifically, the use of an objective experimental paradigm and the interaction of 3-D information with other elements required for interpolation (such as TDs and 2-D relationships) allowed us to rule out a number of possible interpretations not involving 3-D interpolation. The present results are also useful in validating the converging–parallel classification task as a useful tool in the study of 3-D object perception. The task proved to be robustly sensitive to object completion.

Our results leave little doubt that contour interpolation should be viewed as a 3-D process, one that takes as its inputs 3-D positions and orientations of contours and produces as outputs interpolated contours extending through all three spatial dimensions. This conclusion does not merely involve an additional degree of freedom in the standard 2-D formulation of contour interpolation but, rather, suggests considerable changes in the sophistication of the mechanism of interpolation. As stated earlier, these results also place constraints on the neural locus of such a mechanism. The theoretical ramifications of the results are discussed more thoroughly in Kellman et al. (in press).

3-D Relatability in Object Formation

The present results provide support for a notion of 3-D relatability describing the geometry of 3-D interpolation. 3-D relatability is an extension of the theory of relatability, which describes which edges can be connected by interpolation processes in two dimensions (Kellman et al., 2001; Kellman & Shipley, 1991) Object completion in three dimensions appears to be governed by constraints similar to those governing object completion in two dimensions: Interpolated edges connecting across gaps in 3-D space must be smooth (differentiable at least once), monotonic, and in agreement with the orientations of the physically specified edges to which they are connected at their endpoints (points of TD). Below, we consider several particular aspects of this geometry.

Misalignment in depth. Our results show that contour interpolation is disrupted if, from an initially relatable configuration, a depth shift of one contour relative to another results in positions and orientations of the contours that exceed the limits set by 3-D relatability. For interpolation to occur under such circumstances, either the illusory contours would have to doubly inflect, or TDs would need to be introduced. Consistent with 2-D contour interpolation, 3-D contour interpolation appears to be constrained so that interpolation does not occur under these circumstances.

The experiments yielded considerable information about the effects of misalignment. In Experiments 1 and 3, the smaller shift (5.6 arcmin of disparity) showed slightly better performance than the larger shift (11.2 arcmin). Performance levels for all groups in Experiments 3 and 4 (with object formation disrupted for other reasons) was highly consistent with the levels shown in Experiment 1 for the larger shifted displays (despite the different groups of subjects in the experiments). These results suggest that our larger shift eliminated interpolation effects but that our smaller shift did not completely eliminate them. On a metric of disparity shifting, then, relatability is much decreased by shifts on the order of 6 arcmin and eliminated by about 11 arcmin. Putting this in terms of positions in virtual space in the displays we used, relatability was reduced by a 4.5-cm shift (about a 5% shift at a viewing distance of roughly 95 cm) and eliminated by a 9.0-cm shift (about a 9% shift).

There is some question as to the best way to describe the tolerance for misalignment. Relatability specifies boundary conditions under which interpolation does and does not occur, but realistically, small deviations from relatability may be insufficient for completely eliminating the performance benefits in tasks aided by unit formation. Indeed, our results show that the smaller depth shift was less effective in breaking up relatability than was the larger depth shift. In two dimensions, several studies suggest that a shift of 15–20 arcmin of visual angle of initially collinear or cocircular edges (e.g., Palmer, Kellman, & Shipley, 1997; Shipley & Kellman, 1992a) is sufficient to disrupt interpolation. It is likely that this tolerance for misalignment is actually best described not by a fixed retinal value but, rather, by some viewpoint-invariant quantity. In particular, we suspect that the tolerance depends on an angular notion of separation. For example, for two contours oriented vertically, the relevant metric may be the ratio of the horizontal displacement to the vertical separation (i.e., the angle for which this quantity is the tangent). This issue is currently under study for 2-D misalignment in contour interpolation.

The idea that tolerance for deviations from relatability would correspond to an angular metric should also apply to 3-D relatability. In the case of our 3-D stimuli, tolerance would be determined by the ratio of the shift in the horizontal plane to the contour separation. Our experiments indicate that the angle defined by this ratio is rather large—approximately 70°. This estimate, however, does not take into account the possibility that depth misalignment was not veridically perceived. Although the data clearly indicate that both of the shift magnitudes used in the experiments reported here were large enough to allow clear classification of the stimuli as relatable or nonrelatable, determining how best to quantify small deviations from 3-D relatability will have to be addressed in future research.

Constraints on curved interpolation. The present results address the issue of nonrelatability resulting from depth shifts, analogous to the planar shifts that have been shown to disrupt 2-D relatability. Another issue concerns the angular constraints on 3-D interpolation of curved contours. In two dimensions, interpolation seems to become weaker as the angle between the inducing edges becomes more acute (with 180° separation defining collinear inducing edges). When approximately 90° separates the inducing edges, interpolation no longer occurs. Does a similar principle apply to 3-D interpolation? In the present experiments, the most extreme slants did present angles of less than 90°, but again, because of underperception of the stereoscopically specified slants, the perceived angular separation of even the most extremely slanted contours was greater than 90°. We are currently investigating whether a 90° constraint applies to 3-D interpolation.

Torsion. A facet of 3-D interpolation that has no analogue in two dimensions is the concept of surface torsion and how (or if) this is incorporated into the representation of bounding contours. The stimuli used in the present experiments tested relatability in the absence of torsion, but displays can be easily generated to test tolerance for torsion in 3-D contour interpolation. For some point on a space curve parameterized by arclength, $\sigma(s)$, with principle normal vector $\mathbf{N}$ and tangent vector $\mathbf{T}$, torsion, τ, is a scalar that defines how quickly the curve is bending out of the plane defined by $\mathbf{T}$ and $\mathbf{N}$. Torsion can also be defined in terms of the change in the binormal curve, $\mathbf{B}$, and $\mathbf{N}$, at that same point (see Figure 21):

$$\mathbf{B} = \mathbf{T} \times \mathbf{N},$$

and

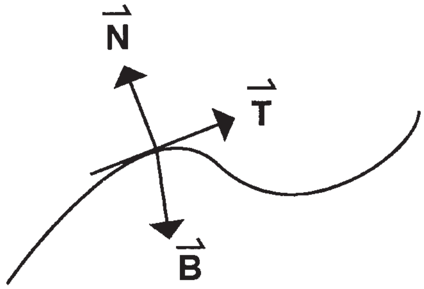

Figure 21. The principle normal (**N**), tangent (**T**), and binormal (**B**) of a space curve.

$$\frac{d\mathbf{B}}{ds} = -\tau\mathbf{N}.$$

Figure 22 shows a series of displays in which unit formation requires interpolated contours with torsion. The amount of torsion necessary increases from top to bottom, and it is apparent that the connectedness of the surfaces correspondingly decreases. In these displays, however, depth shift is confounded with torsion. Figure 23 has a single display in which the visible contours can be connected by a straight interpolated segment and are, therefore, strongly relatable. Here also, however, the surfaces do not appear to be connected. This display suggests that the input into the contour-interpolation process is even more sophisticated than we have suggested above. The lack of interpolation in these displays indicates either that the representation of the contour inputs to interpolation contain information about the surface orientation and that contour torsion prevents interpolation or that contour interpolation is later inhibited by the form of the surface that the interpolated contours bound. Recently work has begun to explore these issues (Fantoni, Gerbino, & Kellman, 2004).

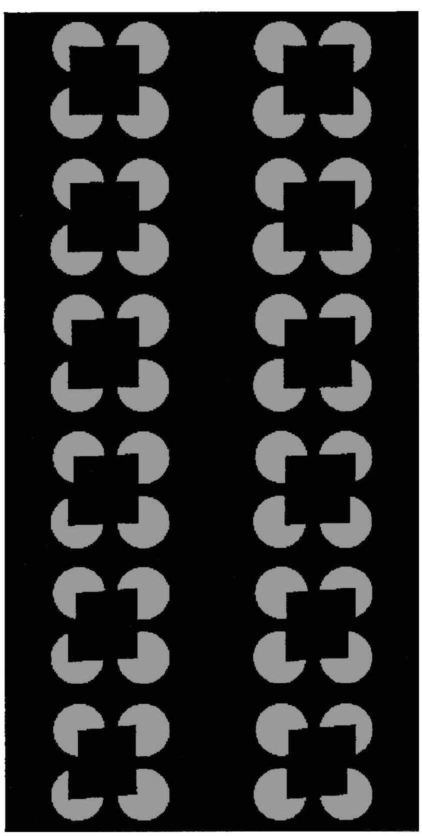

Figure 22. Interpolation displays requiring differing amounts of torsion. The displays are stereo pairs and can be free-fused by crossing the eyes. For interpolation to occur, the illusory contours must have torsion in all but the top display. The amount of torsion necessary for unit formation increases from top to bottom.

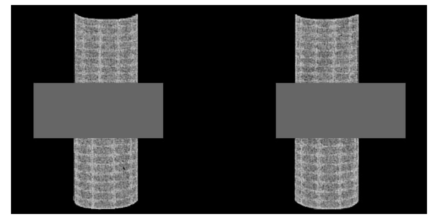

Figure 23. A separate case of torsion? The display is a stereo pair and can be free-fused by crossing the eyes. When the images are fused, the left and right vertical edges of the two half cylinders are collinear. However, the surfaces bounded by the top and bottom edges have different orientations. Contours have positions and orientations in space, but they may also carry an orientation derived from the surface they bound. If contour interpolation depends on all of these properties, these contours cannot be connected without substantial torsion.

Neural Models and 3-D Interpolation

The experiments presented here provide strong support for a 3-D process underlying contour interpolation and the formation and representation of unified objects. Current models of how the visual system extracts the inputs to interpolation (edges and junctions; e.g., Heitger et al., 1992; Morrone & Burr, 1988), contour integration (Field et al., 1993; Yen & Finkel, 1998), and the interpolation process itself (Fantoni & Gerbino, 2003; Grossberg, Mingolla, & Ross, 1997; Heitger et al., 1998) take 2-D inputs modeled after orientation-sensitive units from the earliest visual cortical areas (e.g., V1 and V2). These units are believed to be sensitive to planar orientation—specifically, the orientation of visual sensory input as encoded at the retina. The attempt to tie the inputs used by interpolation processes to known aspects of early visual physiology is laudable. Yet, the evidence regarding 3-D interpolation appears to be incompatible with these models. Interpolation in a plane oriented frontoparallel to the observer may be a special case of more general 3-D interpolation processes. Models that use relations of 2-D orientation-sensitive units cannot be easily extended to incorporate 3-D interpolation.

Reorienting models to include the more general case of 3-D interpolation is far from trivial. Neurons in early visual areas are not sensitive to the stimulus properties necessary to account for the results of the experiments presented here. One might imagine that this problem could be overcome simply by combining their orientation-selective properties with outputs from units sensitive to stereoscopic disparity. This is unlikely to solve the problem for two reasons. First, relative disparity is the quantity necessary for extracting the appropriate depth measurements needed for 3-D interpolation. Relative disparity remains invariant with fixation, whereas absolute disparity changes with the fixation point (being zero at the fixation point, regardless of depth). Cells in V1 have been reported to be sensitive not to relative disparity, but absolute disparity (Cumming & Parker, 1999).

Second, interpolation must incorporate higher level processing because, as we have demonstrated here, it is concerned with properties of contours—such as 3-D position and slant—as they exist in the world (rather than on the retina). Even relative disparities are not capable of conveying depth intervals. Rather, their output must be combined with at least one independent measure of the distance to a point with known relative disparity. Figure 24A illustrates an observer viewing two points separated by a depth interval in the world. The relative disparity given to the observer by this depth interval will decrease with viewing distance. This is the problem of stereoscopic depth constancy (e.g., Wallach & Zuckerman, 1963). The problem becomes more complicated when slant is the quantity that needs to be measured. 3-D slant estimations depend on both the depth interval and the extent over which this depth is traversed (see Figure 24B). The results presented here suggest that the slants of the contours in our stimuli were the relevant quantities for determining whether interpolation will occur. Neurons that code this information are not known to exist in either V1 or V2.

A final consideration for computational and neural models of object formation is that 3-D position in space is known to be specified through a combination of many cues (e.g., stereodisparity, perspective, texture, shading, occlusion). It would seem unlikely that there exist independent 3-D interpolation mechanisms for separate depth cues. Instead, interpolation might occur at some cortical location where 3-D contour positions and orientations are expressed abstractly, independent of the particular cues through which they are specified. These considerations regarding the information and likely cortical loci involved in interpolation may be inconsistent with many recent suggestions (e.g., Bakin, Nakayama, & Gilbert, 2000; Mendola, Dale, Fishl, Liu, & Tootell, 1999; Seghier, Dojat, Delon-Martin, Rubin, & Warnking, 2001; Sugita, 1999) that the earliest cortical areas, V1 and V2, are the places where contour interpolation is accomplished.

Recent neurological evidence suggests that neural units with the appropriate properties may exist in the caudal intraparietal sulcus (cIPS). Sakata, Taira, Kusunoki, Murata, and Tanaka (1997) have located cells in cIPS that are sensitive to 3-D properties of stimuli—in this case, responding to orientation in space specified through stereodisparity or surface texture. Units such as these would be candidates for participation in neural interactions that implement 3-D relatability. Interactions among units encoding 3-D positions and orientations, possibly within a scheme like the "association field" of Field et al. (1993), could account for the results presented here. In short, these neurophysiological findings suggest the existence of units that could be the basis of 3-D interpolation. The key intuition behind some models—that interactions in a network of oriented units could lead to contour interpolation—may yet prove correct. What is much less intuitive from current perspectives is that these interactions may occur in a network whose inputs include 3-D positions and orientations and whose operation computes the 3-D geometry of contour relatability.

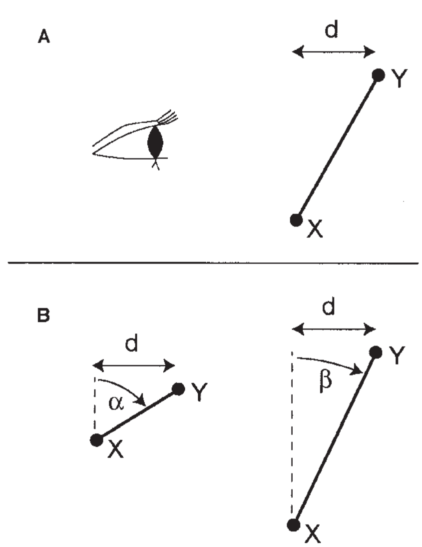

Figure 24. Relations between disparity, edge length, and slant. A: A given depth interval *d* in the world will give rise to decreasing disparity differences between points *X* and *Y* as viewing distance increases. B: Slant depends on both the depth interval between two points and their separation. The two cases shown have the same depth interval *d* between points *X* and *Y*; however, because the separation of *X* and *Y* differs in the two cases, Slant $\alpha >$ Slant β.

References

Albert, M. K. (2001). Surface perception and the generic view principle. *Trends in Cognitive Sciences, 5,* 197–203.

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,* 676–684.

Baylis, G. C., & Driver, J. (1993). Visual attention and objects: Evidence for hierarchical coding of location. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 451–470.

Beard, B., & Ahumada, A. (1998). A technique to extract relevant image features for visual tasks. In B. E. Rogowitz, T. J. Watson, & T. N. Pappas (Eds.), *Human vision and electronic imaging III: SPIE proceedings* (Vol. 3299, pp. 79–85). Bellingham, WA: International Society for Optical Engineering.

Bressan, P., Mingolla, E., Spillman, L., & Watanabe, T. (1997). Neon color spreading: A review. *Perception, 26,* 1353–1366.

Carman, G. J., & Welch, L. (1992, December 10). Three-dimensional illusory contours and surfaces. *Nature, 360,* 585–587.

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.

Day, R. H., & Kasperczyk, R. T. (1983). Amodal completion as a basis for illusory contours. *Perception & Psychophysics, 33,* 355–364.

Dumais, S. T., & Bradley, D. R. (1976). The effects of illumination level and retinal size on the apparent strength of subjective contours. *Perception & Psychophysics, 19,* 339–345.

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

Fantoni, C., Gerbino, W., & Kellman, P. J. (2004). Approximation, torsion, and amodally-unified surfaces [Abstract]. *Journal of Vision, 4*(8), 726a.

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.

Gans, D. (1969). *Transformations and geometries.* New York: Appleton-Century-Crofts.

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.

Gillam, B., & Blackburn, S. G. (1998). Surface separation decreases stereoscopic slant, but a monocular aperture increases it. *Perception, 27,* 1267–1286.

Gillam, B., & Ryan, C. (1992). Perspective, orientation disparity, and anisotropy in stereoscopic slant perception. *Perception, 21,* 427–439.

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., & 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 Neurosciences, 20,* 106–111.

Guttman, S., & Kellman, P. J. (2002). Do spatial factors influence the microgenesis of illusory contours? [Abstract]. *Journal of Vision, 2*(7), 355a.

Guttman, S., Sekuler, A. B., & Kellman, P. J. (2003). Temporal variations in visual completion: A reflection of spatial limits? [Abstract]. *Journal of Vision, 3*(9), 240a.

Hakkinen, J., Liinasuo, M., Kojo, I., & Nyman, G. (1998). Three-dimensionally slanted illusory contours capture stereopsis. *Vision Research, 38,* 3109–3115.

Heider, B., Spillman, L., & Peterhans, E. (2002). Stereoscopic illusory contours—Cortical neuron responses and human perception. *Journal of Cognitive Neuroscience, 14,* 1018–1029.

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.

Hoffman, D. D. (1998). *Visual intelligence.* New York: Norton.

Howard, I. P., & Rogers, B. J. (1995). *Binocular vision and stereopsis.* New York: Oxford University Press.

Kanizsa, G. (1979). *Organization in vision.* New York: Praeger.

Kellman, P. J., Garrigan, P., Kalar, D., & Shipley, T. (2003). Good continuation and relatability: Related but distinct principles [Abstract]. *Journal of Vision, 3*(9), 120a.

Kellman, P. J., Garrigan, P., & Shipley, T. F. (in press). Object interpolation in three dimensions. *Psychological Review.*

Kellman, P. J., Guttman, S. E., & Wickens, T. D. (2001). Geometric and neural models of object perception. In T. F. Shipley & P. J. Kellman (Eds.), *From fragments to objects: Segmentation and grouping in vision* (pp. 183–245). Oxford, England: Elsevier Science.

Kellman, P. J., & Loukides, M. G. (1987). An object perception approach to static and kinetic subjective contours. In S. Petry & G. E. Meyer (Eds.), *The perception of illusory contours* (pp. 151–164). New York: Springer-Verlag.

Kellman, P. J., Machado, L. J., Shipley, T. F., & Li, C. C. (1996). 3-D determinants of object completion. *Investigative Ophthalmology and Visual Science, 37*(Suppl. 3), S685.

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

Kellman, P. J., & Shipley, T. F. (1992). Perceiving objects across gaps in space and time. *Current Directions in Psychological Science, 1,* 193–199.

Kellman, P. J., Yin, C., & Shipley, T. F. (1998). A common mechanism for illusory and occluded object completion. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 859–869.

Liu, Z., Knill, D. C., & Kersten, D. (1995). Object classification for human and ideal observers. *Vision Research, 35,* 549–568.

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: Publications Universitaires de Louvain.

Morrone, M. C., & Burr, D. C. (1988). Feature detection in human vision: A phase-dependent energy model. *Proceedings of the Royal Society of London, Series B, 235,* 221–245.

Nakayama, K., Shimojo, S., & Ramachandran, V. S. (1990). Transparency: Relation to depth, subjective contours, luminance, and neon color spreading. *Perception, 19,* 497–513.

Palmer, E. M., Kellman, P. J., & Shipley, T. F. (1997). Spatiotemporal relatability in dynamic object completion. *Investigative Ophthalmology and Visual Science, 38,* 256.

Palmer, E. M., Kellman, P. J., & Shipley, T. F. (2000). Modal and amodal perception of dynamically occluded objects. *Investigative Ophthalmology and Visual Science, 41,* 439.

Petry, S., & Meyer, G. E. (Eds.). (1987). *The perception of illusory contours.* New York: Springer-Verlag.

Punzel, S., Yonas, A., & Schrater, P. (2001). Grouping and detection of global apparent motion [Abstract]. *Journal of Vision, 1,* 164a.

Ringach, D. L., & Shapley, R. (1996). Spatial and temporal properties of illusory contours and amodal boundary completion. *Vision Research, 36,* 3037–3050.

Rock, I., & Anson, R. (1979). Illusory contours as a solution to a problem. *Perception, 8,* 665–681.

Saidpour, A., Braunstein, M. L., & Hoffman, D. D. (1994). Interpolation across surface discontinuities in structure from motion. *Perception & Psychophysics, 55,* 611–622.

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,* 350–357.

Seghier, M., Dojat, M., Delon-Martin, C., Rubin, C., & Warnking, J. (2001). Moving illusory contours activate primary visual cortex: An fMRI study. *Proceedings of the National Academy of Sciences, 98,* 1907–1911.

Sekuler, A. B., Palmer, S. E., & Flynn, C. (1994). Local and global processes in visual completion. *Psychological Science, 5,* 260–267.

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: Evidence for an identity hypothesis. *Journal of Experimental Psychology: Human Perception and Performance, 18,* 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,* 97–106.

Shipley, T. F., & Kellman, P. J. (Eds.). (2001). *From fragments to objects: Segmentation and grouping in vision.* Oxford, England: Elsevier Science.

Singh, M., & Hoffman, D. D. (1999a). Completing visual contours: The relationship between relatability and minimizing inflections. *Perception & Psychophysics, 61,* 943–951.

Singh, M., & Hoffman, D. D. (1999b). Contour completion and relative depth: Petter's rule and support ratio. *Psychological Science, 10,* 423–428.

Sugita, Y. (1999, September 16). Grouping of image fragments in primary visual cortex. *Nature, 401,* 269–272.

Thornber, K. K., & Williams, L. R. (2000). Characterizing the distribution of completion shapes with corners using a mixture of random processes. *Pattern Recognition, 33,* 1683–1711.

van Ee, R., Banks, M. S., & Backus, B. T. (1999). An analysis of binocular slant contrast. *Perception, 28,* 1121–1145.

Vreven, D., & Welch, L. (2001). The absence of depth constancy in contour stereograms. *Perception, 30,* 693–705.

Wallach, H., & Slaughter, V. (1988). The role of memory in perceiving subjective contours. *Perception & Psychophysics, 43,* 101–106.

Wallach, H., & Zuckerman, C. (1963). The constancy of stereoscopic depth. *American Journal of Psychology, 76,* 404–412.

Wertheimer, M. (1938). Laws of organization in perceptual forms. In W. Ellis (Ed.), *A source book of Gestalt psychology* (pp. 71–88). London: Routledge & Kegan Paul. (Original work published 1923)

Wurtz, R. P., & Lourens, L. (1999). Corner detection in color images through a multiscale combination of end-stopped cortical cells. *Image and Vision Computing, 18,* 531–541.

Yarbus, A. F. (1967). *Eye movements and vision.* New York: Plenum Press.

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.

Yin, C., Kellman, P. J., & Shipley, T. F. (2000). Surface integration influences depth discrimination. *Vision Research, 40,* 1969–1978.

Zemel, R., Behrmann, M., Mozer, M., & Bavelier, D. (2002). Experience-dependent perceptual grouping and object-based attention. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 202–217.

Received June 28, 2004
Revision received September 23, 2004
Accepted November 23, 2004 ■