

Volume Completion

Peter Ulric Tse

Harvard University Vision Sciences Lab

The visual system completes image fragments into larger regions when those fragments are taken to be the visible portions of an occluded object. Kellman and Shipley (1991) argued that this “amodal” completion is based on the way that the contours of image fragments “relate.” Contours relate when their imaginary extensions intersect at an obtuse or right angle. However, it is shown here that contour relatability is neither necessary nor sufficient for completion to take place. Demonstrations that go beyond traditional examples of overlapping flat surfaces reveal that “mergeable” volumes, rather than relatable contours, are the critical elements in completion phenomena. A volume is defined as a 3-D enclosure. Typically, this refers to a surface plus the inside that it encloses. Two volumes are mergeable when their unbounded visible surfaces are relatable or the insides enclosed by those surfaces can completely merge. Two surfaces are relatable when their visible portions can be extended into occluded space along the trajectories defined by their respective curvatures so that they merge into a common surface. A volume-based account of amodal completion subsumes surface completion as a special case and explains examples that neither a contour- nor a surface-based account can explain. © 1999 Academic Press

INTRODUCTION

How is the percept of a 3-D scene constructed from the 2-D retinal image projection? Because many states of the world could project to the same image, the image is inherently ambiguous. To overcome this ambiguity the visual system must make assumptions about the relationship between image information and the structure of the world, and it must infer or “add” information about the world that is “missing” from the image. For decades the study of completion phenomena has been at the center of attempts to understand such constructive processes. Several authors (e.g., Kanizsa, 1955, 1979; Kanizsa & Gerbino, 1982; Michotte et al., 1964) have found evidence that “amodal” completion (of an object behind an occluder) and “modal”

Address reprint requests to Peter Tse at the Max Planck Institute for Biological Cybernetics, 38 Spemannstrasse, Tuebingen, 72076, Germany. E-mail: ptse@wjh.harvard.edu. I thank Patrick Cavanagh, Ken Nakayama, Marc Albert, Bart Anderson, Martha Farah, and the anonymous reviewers for helpful comments on earlier drafts. The author was supported by AASERT F49620-94-0376 and McDonnell-Pew 98-49CNS.






FIG. 1. Traditional example of modal completion (a) and amodal completion (b).

completion (of an object in front of inducers) are not due to cognitive inference, but rather are the automatic outputs of perceptual processing.

A classical example of modal completion is shown in Fig. 1a, and an example of amodal completion is shown in Fig. 1b. Early authors emphasized Gestalt organizational principles (Koffka, 1935), such as global stability, regularity, and simplicity of form, to explain why image fragments complete (e.g., Hochberg & McAlister, 1953). More recently there has been a tendency to conceive of completion in terms of good contour continuation (e.g., Kellman & Shipley, 1991; Wouterlood & Boselie, 1992) or surface completion on a common depth plane (Nakayama et al., 1989, 1995; Nakayama & Shimojo, 1992). This paper will offer demonstrations that imply that “traditional” contour- and surface-based theories of completion are too limited. An account of completion in terms of the linking of surfaces and the merging of the insides that they enclose will be developed that can deal with several novel counterexamples to traditional theories.

TRADITIONAL VIEWS OF COMPLETION

Perhaps because traditional demonstrations of completion like those in Fig. 1 could be constructed by laying flat pieces of paper one on top of the other, many researchers have tended to conceive of completion phenomena in terms of 2-D edges and flat surface relationships. But in the real world, occlusion happens between volumetric objects that are bounded by visible and occluded surfaces. Because we must perceive an object’s spatial extent in order to grasp it or maneuver around it, it is not surprising that visual processes have evolved to complete the volumetric extent of objects based on the shape of their visible surfaces. The aim of the present paper is to extend traditional accounts of completion in terms of contours and flat surfaces to the more natural domain of surfaces that enclose space and self-

occlude as well as occlude. Before doing this it will be necessary to summarize the dominant current theories of completion and offer examples that demonstrate the need for a more general approach.

Traditionally, there has been great interest in describing the local image cues that allow the visual system to determine occlusion relationships. One image cue to occlusion is a contour tangent discontinuity, such as a T-junction (Clowes, 1971; Huffman, 1971; Waltz, 1975; Lowe, 1987; Malik, 1987; Nakayama et al., 1989; Kellman & Shipley, 1991). Even though T-junctions are not necessarily present when occlusion is perceived (e.g., Fig. 24; Tse & Albert, 1998), T-junctions are generically present when one surface occludes another surface separated in depth. Another more global image cue that can be used for completion is the relative orientation of image contours. Kellman and Shipley (1991) formalized the Gestalt law of good continuation (Wertheimer, 1923, 1938; Ullman, 1990) and argued that two edges occluded by a single object would amodally complete when the angle between their intersecting imaginary extensions (that have no reversals of curvature) subtended 90° or more (Fig. 2). Because the occluded bars in the lefthand case of Fig. 2a satisfy these conditions, they “relate” and amodally complete behind the occluder. Since contour extensions do not intersect in the righthand case of Fig. 2a, amodal completion fails to take place. Kellman and Shipley argued that contour relatability was the basis of both modal and amodal completion phenomena.

Kellman and Shipley also discussed instances where completion takes place even when contour extensions subtend less than 90° (e.g., the triangle in Fig. 2c) and are therefore not relatable according to their definition. They suggested that completion in such cases might be due to top-down factors. Boselie and Wouterlood (1992) felt that such cases should not be so easily dismissed and instead developed their own theoretical extension of the Gestalt law of good continuation that could account for the completion of the triangle in a bottom-up manner (Wouterlood and Boselie, 1992). Other authors have challenged local-cue-driven accounts of completion like those of Kellman and Shipley or Wouterlood and Boselie and have tried instead to explain completion in terms of global regularities in the patterns of completing objects (e.g., Buffart, Leeuwenberg, & Restle, 1981; Sekuler, 1994; Sekuler, Palmer, & Flynn, 1994; van Lier, van der Helm, & Leeuwenberg, 1994, 1995). While there has been an ongoing debate about the nature of the completion process, it is fair to say that a view built upon the Gestalt law of good contour continuation has become dominant, even though it is widely acknowledged that global figural factors also influence this basic underlying process.

It has also been argued that amodal and modal completion involve surface completion (Nakayama et al., 1989, 1995; Nakayama & Shimojo, 1992). The key concept relating surface completion to amodal and modal completion is “border ownership” (see also Rubin, 1915, 1958). When there is a border




FIG. 2. Examples of Kellman and Shipley's (1991) notion of contour relatability are shown in (a). Contours are relatable in the left-hand case, and are not relatable in the right-hand case. The imaginary contour extensions are made explicit in (b). Completion of the triangle seems to take place in (c) even though the contours are not relatable. Kellman and Shipley (1991) attributed completion of the triangle to the influence of top-down processing.

in an image separating regions projected from two surfaces separated in depth, that border is projected from the nearer of those surfaces. Assuming no accidental alignment of surface edges, the border between two image regions can only project from one of those surfaces. The region that owns the border is taken to project from a surface that occludes the surface that projects to the region that does not own the border. The occluded surface can then continue behind the occluding surface and link with other occluded surfaces on the same depth plane because its corresponding image region is "unbounded" on the side where it does not own its border.

Thus, there have traditionally been two dominant but related families of views of completion. The “good contour continuation view” was based on detecting local image cues to occlusion, such as T-junctions, and testing for good contour continuation over them. Under Kellman and Shipley’s theory the inputs to the completion process are local junctions, contour tangent discontinuities, and contour orientations, and the outputs are global “units,” such as surfaces or holes. The appeal of this view is that these cues to occlusion are measurable in the image, so that given an image, one can predict whether the visual system will complete disjoint fragments. In contrast, according to the “surface completion view,” the inputs into the completion process are image regions that do (or do not) own their border everywhere, and the outputs are surfaces whose edges and relative depths have been specified. The surface completion view involves completion over internal representations rather than image elements such as contours, because unbounded surfaces must be *inferred* from image cues. They cannot be identified directly in the image.

While completion over surfaces involves operations over a higher level space than image space, surface and contour completion are by no means mutually exclusive. For example, the outputs of a stage that tests for contour reliability within some local window could serve as candidate inputs into a surface completion process that assigns ownership to one side of a contour. Reliable contours with compatible “inside/outside” assignments could then avoid being vetoed as potential surface patches. Local potential surfaces consistent with image cues could then constrain one another globally, and the outputs would be completed surfaces. Indeed, if we consider the hypothetical stage that tests for contour reliability to be an early but integral stage of the surface completion process, the surface completion view subsumes the contour completion view. Of course, this is only one way that contour and surface completion may be related. Several authors have postulated other interactions between contour interpolation and surface formation processes (e.g., Grossberg & Mingolla, 1985; Kellman & Shipley, 1991; Yin et al., 1997).

OVERVIEW

The next two sections describe cases of amodal completion that challenge the dominant contour- and surface-based views. Amodal completion is the focus here because amodal and modal completion probably result from a common completion process (Grossberg & Mingolla, 1985; Kellman & Shipley, 1992; Kellman, Yin, & Shipley, 1998). Amodal completion (i.e., occlusion among opaque surfaces) is also far more common in the real world than modal completion (i.e., perfect camouflage). Several counterexamples to traditional theories will demonstrate that amodal completion can fail in the presence of reliable contours in the image (Figs. 3 and 5), and that

completion can take place in the absence of reliable contours (Figs. 2c, 4a, 6, 7, 13, 18, and 19). Other examples (Figs. 4 and 6–12) will demonstrate that the traditional surface completion account needs to be extended to encompass curved surfaces and self-occluding surfaces. Finally, several examples (Figs. 5, 13, 18–21, and 24) will demonstrate the need for a fundamental revision in the traditional surface-based account. A volume-based theory is required in order to account for the role played in completion processes by the space enclosed by visible surfaces. The term “volume” is used in this paper to mean a 3-D enclosure. For the majority of examples considered here this can be taken to refer to a surface plus the material inside that it encloses. Later examples will, however, raise the possibility of volumes that lack distinct surfaces (Figs. 20 and 21).

Thus, the present account of completion first extends the surface-based account of Nakayama and colleagues by emphasizing the way that surfaces interact and “close up” in occluded and self-occluded space. It then goes beyond a surface-based account by emphasizing the role played by a volume’s inside in completion processes. Just as the surface-based account of completion subsumes a contour-based account by assigning an edge an “inside” that belongs to a surface, the present volume-based account subsumes surface-based accounts by assigning surfaces an inside that belongs to a volume. Surface completion will be subsumed under a volume completion account as a special case where completion takes place among “degenerate” volumes (i.e., surfaces) that do not have insides. A unified account of amodal completion will emerge that describes surface merging and the merging of insides as two aspects of a unitary volume completion process.

COUNTEREXAMPLES TO THE CONTOUR RELIABILITY VIEW

Although the visual system appears to use the reliability of contours as one cue to completion among others, there are examples (Fig. 3) where reliable contours in the image do not lead to a completion percept. These examples demonstrate that 3-D surface orientation relationships and depth placement may veto unit formation despite the reliability of contours in the 2-D image (compare Kellman & Shipley, 1991, Fig. 30). A more precise explanation of why completion fails here will be given later.

Conversely, for objects that self-occlude, objects can complete in the absence of reliable contours in the image. For example, in Fig. 4a, the visible portions of the object complete in the absence of visible contours that could link up by simply extending those contours. In order to complete occluded objects, the visual system cannot rely solely on linking visible contours, because the contours that presumably link can themselves be occluded.

Even when there are contours that might join up behind an occluder, as in Fig. 4b, there is no need for the angle to be 90° or greater, as Kellman and Shipley (1991, p. 175) maintained (cf. Boselie & Wouterlood, 1992).




FIG. 3. Amodal completion fails to take place in spite of relatable contours in (a), (b), and (c).

Indeed, the angle between the inner contours defining the two occluded “legs” of this figure (indicated by arrows) define an acute angle, and still the legs seem to join smoothly behind the occluder.

Thus, we have seen examples of relatable contours without amodal completion (Fig. 3) and amodal completion without relatable contours (Fig. 4). *Contour relatability is therefore neither necessary nor sufficient for amodal completion to occur.*




FIG. 4. (a) Amodal completion takes place in the absence of reliable contours. (b) Contours relate even though their imaginary extensions subtend an acute angle.

COUNTEREXAMPLES TO THE TRADITIONAL SURFACE COMPLETION VIEW

Figures 5a and 5b are not only counterexamples to the contour reliability account of amodal completion, they are also counterexamples to the traditional surface completion view. According to this view unbounded surfaces on a common depth plane will join behind an occluder. In Fig. 5a two distinct blocks are perceived behind the occluder even though there are two faces of the left block that are coplanar with faces of the right block. According to the traditional surface completion view these faces should complete. But some of the surfaces that might link in Fig. 5a are not visible. Perhaps completion must take place over visible surfaces. In Fig. 5b, however, this is not a problem. Two visible surfaces of the cubes lie on a common depth plane yet do not complete. Completion may fail in these examples because the surfaces appear to belong to two distinct volumes, and a single surface can only belong to a single volume at a time. Or completion may fail here because the inside of one object's surface cannot link with the outside of another object's surface to form a single topologically closed surface. This and other examples will suggest that the inside/outside assignments of two or more surfaces must be consistent in order for completion to take place. But once we start assigning a role for inside/outside in amodal completion, we begin to leave the domain of pure surfaces (which have no inside or outside) and enter the realm of volumes.




FIG. 5. Amodal completion fails in (a) and (b) in the presence of both reliable contours and reliable surfaces.

SURFACE RELATABILITY

This section will consider examples where completion must take place over inferred entities, such as surfaces or volumes, rather than image entities, such as contours. Since there is a tradition of thinking about completion in terms of surfaces, the aim of the next two sections is to describe as much about completion in terms of surfaces as possible, without reference to volumes.

The notion that unbounded surfaces join in occluded space can be extended from the traditional flat surface domain to the curved and self-occluding surface domain using the concept of surface relatability. *Two surfaces are relatable when their visible portions can be extended into occluded or self-occluded space along the trajectories defined by their respective curvatures so that they smoothly merge into a common surface.* This intuitive




FIG. 6. A tube completes in the absence of any occluding contour projecting from the inferred tube.

definition can be formalized in various ways.¹ But this definition is really a “straw man” because examples of amodal completion in the absence of surface relatability will be given later (Fig. 13) to show that a purely surface-based account is inadequate and that completion must be analyzed at a level of mergeable volumes. For the time being, however, let us see how far the notion of relatable surfaces will get us in our attempts to understand amodal completion.

Before we reject this or any other notion of surface relatability as an insufficient account of completion, let us see what it can explain that cannot be explained in terms of contour relatability or the linking of flat surfaces. Immediately, we can say that completion fails in Fig. 3 because the surfaces on either side of the occluders are not relatable. Surface relatability could also explain why completion occurs in Fig. 4. Next, consider Fig. 6. Although

¹ For those readers who would like a more formal definition of surface relatability, a useful one that reduces surface relatability to the familiar concept of contour relatability might be the following. Consider a plane P that intersects two smooth visible surface patches $S1$ and $S2$ in the neighborhoods where they become unbounded because of the interposition of an occluder or because of self-occlusion. The intersection of P with $S1$ and $S2$ will be two visible unbounded space curve segments $C1$ and $C2$ lying in $S1$ and $S2$, respectively. (Do not consider cases where the intersection forms loops.) $C1$ and $C2$ will have arbitrarily short lengths determined by the size of the patches under consideration. For every pair of $C1$ and $C2$ consider a third space curve $C3$ in occluded space also lying in P . The visible surface patches $S1$ and $S2$ are relatable if, for all P , a $C3$ exists such that (1) $C3$ has no reversals of curvature, (2) $C3$ is everywhere differentiable, and (3) $C3$ extends $C1$ and $C2$ along their curved or straight trajectories until they meet such that the first and second derivatives at the two points where $C3$ meets $C2$ and $C1$ are well defined for all $C1$, $C2$, and $C3$. Since C , the union of $C1$, $C2$, and $C3$, is a smooth space curve lying in P , its image projection C' , will be a smooth contour, and $C1'$ and $C2'$ will be relatable according to some appropriate definition of contour relatability, such as that of Takeichi et al. (1995), if $S1$ and $S2$ are relatable.




FIG. 7. How did your visual system infer a closed surface which is the shape of a spiral?

there is no explicit rendering of a tube here, most observers report seeing a tube connecting and merged with two “squares,” creating a hollow “dumb-bell.” This example is difficult for the traditional surface completion view to account for because a tube completes in the absence of any occluding contour in the image projected from the thin tubular portion itself. Moreover, the front square owns its outer border everywhere and, according to the traditional view, would not be unbounded. However, it *is* unbounded because it contains a self-occlusion in the hole, and it is this inner surface that is relatable to the underside of the occluded square.

Another example is shown in Fig. 7. The perceived shape is typically regarded as a spiral “seashell” by naive observers, even though there is no explicit rendering of a spiral in this image. Traditional surface completion procedures based on relating contours behind an occluder, or linking unbounded surface fragments into a larger surface on the same plane as those fragments, are not sufficient to generate a 3-D percept of a spiral. In particular, note that there are no relatable visible contours between segments of the spiral. Significant local cues to 3-D form, such as the curled contour endpoints of the front portion (cf. Koenderink & van Doorn, 1982) may propagate form information to distal parts of an image that lack such information. Moreover, the shape of the occluding contours seems to indicate that the surfaces curve back on themselves in a way consistent with a spiral interpretation.

Note, however, that if the front portion of Fig. 7 is altered to look like a




FIG. 8. Change the front portion of figure 7, and a spiral is no longer perceived.

bounded surface, as in Fig. 8, a spiral cannot complete, because now there is no spiral form information to propagate across the remaining ambiguous portions of the image. In this case, the semioval fragments (which have not been altered) no longer appear to comprise the loops of a spiral. They may be discs or “eggs,” and they may or may not be contiguous with one another in 3-D. In the next section a surface relatability account of how and why surfaces close up on themselves in self-occluded space to form spirals or eggs will be given.

SURFACE CLOSURE AND CONNECTEDNESS

According to the above italicized definition of surface relatability, not only can surfaces link up behind an occluder, the visible surface regions of an object can relate to one another and close up in self-occluded space (behind the object’s visible surface) to form a volume. For example, when we look at a ball, we do not perceive a hemisphere, we perceive a sphere. Why is this, when the image could be the projection of a hemisphere? A simple argument can be made using the “nonaccidental viewpoint assumption” (Binford, 1981; Barrow & Tenenbaum, 1981; Rock, 1984; Lowe, 1985; Richards et al., 1987; Nakayama & Shimojo, 1992). The visual system operates as if it made the assumption that it is not perceiving an object from one of the few “accidental” viewpoints from which an object’s surface layout is not derivable from its projected contours and other image cues. In the case of the sphere, since there is no image evidence that the visible surfaces do not continue into self-occluded space, it is assumed that they do. The points where the line of sight grazes the surface (the “rim”; Koenderink, 1990) are not taken to be points where the surface abruptly ends. Instead, such




FIG. 9. (a) Although the image projection is 2-D like a disc, the occluding contour may correspond to the rim of a smooth surface. (b) The line of sight grazes the surface at the rim and lies in a surface tangent to the presumed surface. (c) The smoothest surface that links these inferred surface tangents is generated for both the front and the back of the object.

points are seen to be places where the surface begins to occlude itself and continue into self-occluded space. One implication of the nonaccidental viewpoint assumption, then, is that visible surfaces will complete on their far side, and close into a volume, if visible surface curvatures are such that their extensions into occluded space would close (Fig. 9).

How precise is the interpolated form of the occluded or self-occluded portion of an object? Several studies have been concerned with the recovery of the shape of the occluded contour of a partially occluded object based on visible contours' extensions into occluded space (Gerbino & Salmaso, 1985, 1987; Sekuler & Palmer, 1992; van Lier et al., 1994, 1995; Takeichi et al., 1995; Takeichi, 1995). When the contours that complete are themselves occluded, the visual system seems uncertain of the shape of the occluded portions of the object. Most observers agree that they can imagine seeing the occluded region in Fig. 4a as either "fat" or "skinny," as shown in Fig. 10, but that certain occluded forms are more probable than others. This suggests that the visual system may determine *that* amodal completion has taken place without necessarily interpolating the exact form of the occluded region.




FIG. 10. The occluded form of the object in Fig. 4a is open to interpretation. It may be “skinny” as in (a) or “fat” as in (b).

To the extent that an occluded form is interpolated, it may be probabilistic in nature.

Even though occluded form may not be precisely known, we are generally quite good at describing the basic structure of a volume. For example, observers generally report seeing the “worm” in Fig. 11 to be “tubular” like a sausage. What information in the contours of this figure causes the visual system to generate a tubular form when there are no other cues to tubularity such as shading, highlights, or texture gradients? One constraint on possible



FIG. 11. Why does your visual system see this “worm” as approximately tubular in shape?




FIG. 12. (a) Waves receding into the distance or overlapping flat surfaces? (b) A phenomenally closed surface in the absence of a bounding contour. (c) Two overlapping flat surfaces or a single cylindrical “polyp”?

solutions could be the assumption that the surface that projects no visible contours from its interior is smooth. This follows because not only the presence of image contours, but also a *lack* of contours is assumed to be nonaccidental (Grimson, 1981; Brady & Grimson, 1981; Witkin and Tenenbaum, 1983; Barrow & Tenenbaum, 1981; Richards et al., 1987). A lack of image contours means that there is nothing in the projecting surface, such as folds or gaps, that would have given rise to such information in the image. This smoothness assumption may provide a plausible explanation for the perception of tubular form and may account for why the far side of the worm is presumed to be closed. This is because the tangent plane at the rim is assumed to lie along the line of sight, and the smoothest surface with most uniform curvature that connects the rim via those tangent planes is curved on both the near and far sides.

The smoothness assumption also implies surface connectedness, even when surface closure is not possible. In Fig. 12a (Koffka, 1935), the successive layers do not seem to be overlapping, flat, disconnected surface layers. Rather, they seem to most observers to be smoothly connected waves or mountains receding into the distance. In Fig. 12b, there is no occluding contour specifying a closed surface. Yet the disjoint unbounded surfaces speci-

fied by several unclosed occluding contours seem to link into a single lumpy closed surface. Similarly, in the case of the “polyp” shown in Fig. 12c, observers do not report one flat surface occluding another disconnected flat surface. Rather, the near and far surfaces of the polyp seem to link into a connected cylindrical shape.

Since there are no explicit cues that these surfaces connect, they connect because there is no image evidence to veto their connection. In the absence of image evidence to the contrary, unbounded surfaces will link up in occluded or self-occluded space, making surface connectedness a default of visual processing. With regard to completion, *image contours may be less cues of permission than cues of denial*. Only when image contours indicate nonrelatable surfaces will unbounded surfaces not link up. Thus, a consequence of relating surfaces in occluded and self-occluded space appears to be a maximal closing of surfaces. This is not only because an unbounded surface will link with other unbounded surfaces until such linkages are exhausted (Nakayama & Shimojo, 1992), it is also because unbounded surfaces can close up on themselves to become volumes.

VOLUME MERGEABILITY

The previous two sections described examples of amodal completion that could not be accounted for in terms of the linking of relatable contours or surfaces on a common depth plane. However, they could be accounted for in terms of the linking of relatable curved surfaces that self-occlude. This section describes examples where even an account of completion in terms of surface relatability is inadequate, and an analysis in terms of the merging of the insides enclosed by visible surfaces is required.

While surface relatability, like contour relatability, may play an important role in completion, *surface relatability is neither necessary nor sufficient for completion to take place*. Amodal completion fails in Fig. 5b in the presence of both relatable contours and relatable surfaces. Conversely, completion takes place for the five examples shown in Fig. 13 in the absence of relatable contours or relatable surfaces. It seems that amodal completion takes place in the examples shown in Fig. 13 because the insides or substance of the visible portions of the respective objects can merge. For example, consider Fig. 13b. Amodal completion may take place here because the “gel” lying on top of the occluder seems to pass through the hole. This suggests that amodal completion may involve the merging of the insides of volumes.

The type of information over which completion takes place in the examples shown in Fig. 13 may be one of “unbounded volumes.” An unbounded volume is an unbounded surface plus its unbounded inside that could have generated the various image cues to surface layout and depth placement. Bounded volumes, in contrast, are closed surfaces plus the insides they enclose.




FIG. 13. Amodal completion takes place in the absence of reliable contours or reliable surfaces. Completion may take place because insides merge. (a) and (c) The black portions seem to belong to a single object penetrating the “slabs.” (b) The gel appears to pass through the hole. (d) The blob appears to pass around the back of the occluder. (e) A unitary blob seems to pass through a tight “collar.”

Volume “mergeability” means that the inside of one unbounded volume can join the inside of another unbounded volume to create a larger unbounded or bounded volume. How insides merge will be considered later. In cases like those shown in Fig. 13, amodal completion takes place because insides can merge behind (or through) an occluder, even though visible surfaces are not reliable.

Even the central tenet of the traditional flat surface approach, namely, that




FIG. 14. The “arms” appear to wrap around the cylinder whether they are volumetric, like the top three, or volumeless, like the lowest example. Note that the amodally completing portions own their borders with the cylinder.

occluders own all their borders, need not be true because amodal completion is determined only at a level of representation where surfaces and the insides that they enclose can merge and wrap around an occluder. In the traditional domain of flat surfaces separated by a depth discontinuity, the image projection of the occluder must own all its borders with the image projection of a surface that it occludes. However, volumes can be mergeable even when partially occluded volumes own all their borders with the occluder. For example, in Fig. 14, apart from an arbitrarily small portion of occluding contour that belongs to the cylinder, the gray “worms” own all their borders with the occluding cylinder. A flat ribbon such as the lowest “worm” wrapping around the cylinder indeed does own all its borders with the cylinder. Similarly, the white speckled partially occluded volumes in Fig. 15 amodally complete behind or around the black cylinder even though they own all their borders with the cylinder. This should make clear that *contours are not direct cues to occlusion as such, but rather they are cues to volume formation. Border ownership, far from being a bottom-up cue to amodal completion, is itself only specified once occlusion relationships among volumes have been determined.*

COMPUTING VOLUME MERGEABILITY

How do unbounded insides merge? In order to merge insides, insides must be specified. Past work on the perception of inside/outside relations has




FIG. 15. Traditionally it was thought that an occluder owns all borders with that which it occludes. Here, the occluded volume owns its borders everywhere with the cylinder that partially occludes it.

tended to focus on points within a contour lying on a plane (e.g., Sutherland, 1968; Kovacs & Julesz, 1994; Ullman, 1996). How might the visual system determine the inner space enclosed by a visible curved surface? When considering closed surfaces, “inside” and “outside” are well-defined terms. However, if we consider unbounded volumes, such as a mountain emerging from the earth, the meaning of “inside” becomes more ambiguous. A point just beneath the summit of a mountain such as the point z shown in Fig. 16 would probably count as being inside the mountain, but a point deep in the ground below the mountain’s base, such as point p , might not, even though it would still count as inside the earth. There are many ways to formalize our intuitive understanding of 3-D insidiness.² The essence of that intuition

² Three possible formalizations might be (1) a point is enclosed by a surface for which at least one plane exists that contains the point and intersects the surface in a closed loop that surrounds the point; (2) a point is contained by a surface if half or more of the rays originating from that point intersect the surface; (3) all points on a chord that begins at a visible point and ends at a visible point of an opaque surface without penetrating that surface such that the chord itself becomes visible lie inside that surface.




FIG. 16. The dark line depicts a vertical cross-section of a mountain. The point z lies just below the summit and the point p lies deep in the earth below the summit. Note that z appears to lie inside the mountain, but p does not.

is that points surrounded by an overall concave surface are enclosed by that surface.

However the visual system specifies insideness, a central problem concerns when in the visual processing stream the inside of a surface is specified. The approximate direction of inside can be specified early, without even first specifying inward or outward surface normals (i.e., surface orientation), because the inside of an opaque volume always lies farther away in depth than its visible surface. In other words, the inside of an opaque volume always lies behind or under its visible surfaces.

Now that we have briefly explored how insides might be specified, we can address the way that unbounded insides might merge. A rough but useful metaphor for the inside contained within a surface is expanding “foam” held within a container under pressure. When the surface is closed there can be no “expansion” of the inside out of the surface. But when the surface is unbounded (i.e., passes behind an occluder), the inside that it encloses is free to grow out from the surface in an attempt to link up with available unbounded insides of other unbounded volumes. The flow of an unbounded inside is to some extent constrained by the orientations and curvatures of the bounding visible surface as it passes behind the occluder. The reason insides are constrained to a trajectory specified by visible surfaces is that insides are typically contained by surfaces, and surfaces relate along trajectories defined by their visible curvatures. In the degenerate case of flat surfaces, such as shown in Figs. 1 and 2, the flow of the surface’s contour-bound internal region is determined by the orientation and curvature of the contours heading under the occluder, providing overlap with aspects of traditional theories built upon the Gestalt law of good contour continuation (Kellman & Shipley, 1991; Wouterlood & Boselie, 1992; Takeichi et al., 1995). A 3-D example of the flow of the unbounded inside of the seashell depicted in Fig. 7 is shown in Fig. 17a, and flow for the 2-D case shown in Fig. 2c is depicted in Fig. 17b.

But surface or contour relatability is not the only determining factor of volume mergeability. Even when visible contours or surfaces cannot relate, as in Fig. 13, insides can merge to permit amodal completion, because insides




FIG. 17. The unbounded inside will orient along visible surfaces as it “expands” into occluded space because the inside is contained by a surface, and surfaces relate based on their visible curvature. A 3-D example of the “flow” of inside is shown in (a) and a “degenerate” 2-D example where the interior is contained by visible contours is shown in (b).

can flow out into occluded space wherever a volume becomes unbounded by an occluder. In Figs. 13a–c, and e, the insides merge because they flow through another volume that is perceived to have a hole in it. In Figs. 13d, 14, and 15 the inside flows around the occluding block or pole even though the occluded volume owns all its borders with the occluder.

COMPLETE MERGEABILITY

Note that the examples in Figs. 18a and 18b have unrelatable contours and unrelatable implied surfaces. However, example 18a seems to complete,



FIG. 18. The “cue” of complete mergeability. (a): The unbounded interior of the left-hand rectangle completely merges with the right, and can therefore amodally complete. (b): Amodal completion fails because complete mergeability fails to both the right and left.




FIG. 19. If you see these two cylinders as merging, is there a surface tangent discontinuity where they merge or not?

whereas example 18b does not. This can be accounted for by a requirement for “complete mergeability”: When all the unbounded inside of one unbounded volume can merge with the unbounded inside of another, observers should see those two unbounded volumes as merged (Tse, 1999). Mergeability cannot be measured in the image the way contour relatability can, since volumes must be inferred from image cues. Note that the requirement of complete mergeability only has to be met by one of two amodally completing unbounded volumes. For example, volume completion takes place in Fig. 18a because the unbounded inside of the lefthand “peg” can fully merge with that of the larger peg on the righthand side, even though all of the unbounded inside of the righthand peg cannot fully merge with that of the lefthand peg. In Fig. 18b, however, merging is partial both ways, so completion also fails. Note that the requirement of complete mergeability also holds for the degenerate case where flat unbounded surfaces, rather than unbounded volumes, merge. Amodal completion seems to be strongest for cases where mergeability is complete both ways, as in Fig. 1b.

Unlike the requirements imposed by surface relatability, complete mergeability permits completion even when the form completed behind the occluder would have surface (or, in the degenerate case of flat surfaces, contour) tangent discontinuities. Under this conception of completion, completion takes place over the space or “stuff” contained by surfaces. Interiors are guided into occluded space by partially occluded surfaces (or contours in the degenerate case). As long as the insides can merge, it does not matter if the interpolated surfaces (or contours) are not smooth or precisely

specifiable. This indifference to the precise form of occluded surfaces (or contours) is made evident in Fig. 19. Most observers who see these cylinders as merging are not quite sure whether they merge smoothly (so that together they would have the shape of a wine bottle) or whether they merge at a surface tangent discontinuity (so that together they would resemble the back of a car muffler). Thus, as stated with regard to Fig. 4a and 10, the visual system may note *that* completion occurs without bothering to precisely specify the layout of occluded surfaces (or contours).

SURFACELESS VOLUMES

In this paper we have so far focused on volumes such as rocks that are material enclosures bounded by the surface of that material. However, in examples such as fogbanks, pillars of sunlight through dusty air, swirls of silt, or jets of flame, volumes can complete in the absence of well-defined surfaces. For example, in Fig. 20, the “snake” might be made of smoke. Even though surfaces here are neither opaque nor spatially well specified, completion of a (nonopaque and spatially indistinct) volume takes place. It seems that what is completing here is the space filled by the unbounded volumes rather than completion of precise surfaces.

Another example is the “surfaceless volume” spanning the space occupied by a group of objects. For example, the flock specified by many starlings in flight can be perceived to have an inside as well as an outside. The question of how the visual system groups is an old problem whose best “solution” to date seems to be the various descriptive Gestalt “laws,” such as common fate and proximity (Koffka, 1935). However groups are specified, completion may operate over the surfaceless volumes of space occupied by the group. Thus, the droplets in Fig. 21a appear to group together through the black tube-like object. Similarly, the swarm of “locusts” behind a tree trunk in Fig. 21b should be perceived as one swarm rather than two. Note, that



FIG. 20. Completion in the absence of distinct surfaces.




FIG. 21. (a) Why do the droplets seem to amodally complete through this “tunnel”?
(b) This swarm of “locusts” amodally completes in the absence of any completing surface.

pattern completion may also play a role in amodal completion here (Tse, 1999). It seems that unbounded occupied space can complete with distal unbounded occupied space, whether this is the surface-enclosed material inside an object (e.g., a rock), the surface-enclosed space outside of an object's material (e.g., a room), or the surfaceless space spanned by a group of objects (e.g., a flock). Although in earlier sections we equated the meaning of "volume" with surface-bound enclosure, we can now see that "volume" refers more generally to occupied, filled, or enclosed space, whether bounded by a distinct surface or not.

OTHER FACTORS CONTRIBUTING TO VOLUME COMPLETION

This paper has described a bottom-up account of volume completion. However, there is certainly more contributing to volume completion than surface relatability or volume mergeability. World knowledge may also influence the completion process (Tse, 1998). For example, when viewing a building from a curb corner, only two surfaces of the building are visible. How do we know that the building is cubic? In Fig. 22a, suggested by Marc




FIG. 22. (a) This appears to be a volume like a building, whereas (b) does not.




FIG. 23. If the occluder is interpreted to be a piece of cloth, the shape of the occluded object emerges even when none of its surfaces are visible, as on the left.

Albert, most people report seeing a cubic volume, like a building, behind the occluder. This is surprising because the occluded figure might have been projected from a folded piece of paper or any number of noncubic objects, such as the one shown in Fig. 22b. Presumably, the bias to see volume in such cases is due to object knowledge or templates that inform our judgments about likely 3-D form (cf. Tse, 1998). There may even be an assumption that the back of an object is like the front.

Also, the perception of substance may play a role in completion. For example, in Fig. 13b, the two halves may complete in part because both halves appear to be comprised of the same stuff (cf. Yin, Kellman, & Shipley, 1997). Similarly, if the visible material in Fig. 23 is assumed to be cloth, then a “donut” and half occluded “smoker’s pipe” may seem to complete under the cloth and support it. When and how information about substance constrains completion processes is an open and important question.

GENERAL DISCUSSION

This paper has demonstrated that traditional approaches to amodal completion based on contour relatability and the linking of unbounded surfaces over a common depth plane are too limited. The surface completion view of Nakayama and colleagues was extended to the realm of curved and self-occluding surfaces. But because of the importance of the merging of insides even this elaborated surface completion view was seen to be inadequate. A volume completion view was developed that can account for all the novel demonstrations considered here.

Amodal completion does not only happen behind an occluder. It is a universal aspect of volume completion, since all objects self-occlude their far side and therefore occlude their true extent. The real problem that has been investigated here, then, is not amodal completion at all, but 3-D shape formation or volume completion. What has traditionally been called “amodal com-

pletion'' (or ''modal completion''; see Tse, 1998) is just a small subset of all volume completion phenomena.

What is the nature of the volume completion process? Central questions addressed in this paper were the following: (1) What are the inputs to the completion process? and (2) what are the outputs of this process? According to the account of completion developed here, the inputs are (1) unbounded local surface patches plus their local insides, and the outputs are (2) maximally closed surfaces plus the insides that they specify. The inputs can be thought of as local ''potential volumes'' consistent with image cues. The completion process may involve a relaxation into a best-fit given all these local constraints. Because we are talking about a surface plus the inside that it specifies at both the input and output levels, this is a process of volume completion.

Two new types of information seem to play a role in volume completion: surface relatability and volume mergeability. According to surface relatability, visible surfaces continue into occluded or self-occluded space along the trajectories defined by their visible curvatures and link up if those surface extensions meet smoothly. According to volume mergeability, not only do visible surfaces link and close up maximally, the insides enclosed by those surfaces merge. Are surface-relating and inside-merging separate processes that operate over independent types of information or are they aspects of a single completion process that operates over a common type of information? Does surface-relating happen first, followed by a stage of inside-merging or do surface-relating and inside-merging happen together? The next several paragraphs attempt to answer these questions.

Although they are mathematically different, the distinction between a surface and a volume is a false dichotomy when talking about the natural world. An object always has surfaces, and all natural surfaces have an inside because they are the outermost portions of the materials of things. Even a very thin object like a leaf, or a perfectly flat natural surface like the ground plane, will have an inside, which is whatever is just under the visible surface. However, although not independent in the world, it may be that surfaces and volumes are independent types of representations in the brain. For example, the representation of a ''thin'' maple leaf might be of the ''volumeless'' sort, while the representation of a ''fat'' cactus may be of the volumetric sort. This paper offers no evidence to settle whether there are one or two types of representation. There might, for example, be a stage of surface completion whose outputs are volumeless surfaces, followed by (or operating in parallel with) a stage of inside interpolation and volume merging. Or, there might be a single process of volume completion.

A good case can be made that there is a single process of volume completion, even though it is difficult to conceive of a crucial experiment that might settle the matter. There are no examples that surface completion can explain which a volume completion account cannot explain. But the reverse is true.

Figures 5, 13, 17, 18, and 19 are examples in which a volume completion account succeeds but a surface completion account fails. Since both surfaces and volumes are entities inferred from the image, nothing useful is gained by positing two distinct types of representation. In the absence of evidence that there are two types of internal representation, it is most economical to assume that there is only one. This is especially so because the representation of a volume (a surface plus its inside) subsumes the representation of a surface.

Moreover, volume mergeability cannot be analyzed by the visual system independently of an analysis of surface layout because insides are specified by surfaces. Conversely, an analysis of surface layout will provide certain information about the direction of inside at no computational cost, since the inside of a material always lies behind or under its bounding opaque visible surface. It is therefore likely that at least approximate surface orientation and the direction of inside are specified together. If that is so, then the relating of unbounded surfaces may entail the merging of unbounded insides and vice versa. It seems plausible, then, that surface relatability and volume mergeability are really aspects of a single process of volume completion. While the outputs of the completion process may appear thin or fat, they are the outputs of a common constructive process and thereby comprise a unitary type of representation.

Thus, it seems that volume completion subsumes traditional surface completion accounts. Volumes can complete in the absence of relatable surfaces (Figs. 13 and 18) and fail to complete in the presence of relatable surfaces (Fig. 5). However, relatable surfaces that enclose insides must *always* enclose mergeable insides. If insides are not mergeable, surfaces will fail to complete even if they are relatable. Only relatable surfaces that do not enclose an inside can complete in the absence of volume mergeability. This kind of “volumeless” surface completion is a degenerate case of volume completion and one not found in nature.

Similarly, surface ownership subsumes the notion of edge ownership as a special case. Determining the direction of a volume’s inside can be thought of as a 3-D analog of the determination of edge ownership (or, in the image, border ownership). Only now, occluding contours correspond to oriented surface patches (at the rim) rather than edges, and the direction of inside specifies the volume that owns the surface rather than the surface that owns the edge.

Finally, what is the relationship of volume completion to the lower level types of completion described, for example, by Rensink and Enns (1995, 1998)? The inputs to the volume completion process must themselves be outputs of earlier completion processes that link co-oriented edges into contours and group adjacent regions with similar image properties together. While self-occlusion is not a problem at the early stage of contour completion, noise and local occlusion are, so we should expect some form of amodal



FIG. 24. In the “seamonster” there is amodal completion in the absence of T-junctions or other local cues to occlusion. There is not even an explicit occluder here to allow the surfaces to become unbounded.

completion to operate quite early. Indeed, it is likely that both local and more global completion processes have aspects in common, such as a sensitivity to co-oriented contours. However, these completion processes must be distinct.

No early completion process can account for completion for a case like that shown in Fig. 24 (Tse, 1998; Tse & Albert, 1998). Here there are no local image cues, such as T-junctions, that could indicate occlusion. There is no good continuation or “reliability” between contours that link (Kellman & Shipley, 1991) and there is not even an explicit occluding surface to create unbounded surfaces as required by the traditional surface completion account. Only at a higher level where the individual black portions are processed as potential volumes could the parts complete into the volumes that we perceive.

Early completion processes presumably culminate in the generation of local potential volumes on the basis of contour and other image information, such as texture and luminance gradients. The volume completion process takes these potential volumes as inputs and generates volumes as outputs. We can imagine several interacting stages of completion each of which generates candidates that get vetoed or accepted by the next level based on cues of denial or permission. The ultimate purpose of this cascade of completion processes would be to generate a percept of 3-D form despite incomplete, noisy, and ambiguous information at the retina. On this account, volume completion would be at the highest level of the completion cascade. The outputs of volume completion could then be used as inputs by, for example, object recognition processes that require a match of form to memory or motor processes involved in grasping or maneuvering around the 3-D environment. As such volumes are quite likely the central representational format of mid-level vision.

CONCLUSIONS

Relatable contours and surfaces are neither necessary nor sufficient for completion to take place. Therefore an account of completion based on the linking of visible contours or surfaces behind an occluder is not adequate. A volume completion account was developed according to which completion takes place both over surfaces and the space that they enclose. A volume-based account of amodal completion subsumes surface completion as a special case and explains the foregoing examples that neither a contour- nor a surface-based account can explain.

REFERENCES

- Barrow, H. G., & Tenenbaum, J. M. (1981). Interpreting line drawings as three-dimensional surfaces. *Artificial Intelligence*, **17**, 75–116.
- Brady, M., & Grimson, E. L. (1981). *The perception of subjective surfaces*. A.I. Memo No. 666, MIT.
- Binford, T. O. (1981). Inferring surfaces from images. *Artificial Intelligence*, **17**, 205–244.
- Boselie, F., & Wouterlood, D. (1992). A critical discussion of Kellman and Shipley's (1991) theory of occlusion phenomena. *Psychological Research/Psychologische Forschung*, **54**, 278–285.
- Buffart, H., Leeuwenberg, E., & Restle, F. (1981). Coding theory of visual pattern completion. *Journal of Experimental Psychology: Human Perception and Performance*, **7**, 241–274.
- Clowes, M. B. (1971). On seeing things. *Artificial Intelligence*, **17**, 79–116.
- Epstein, R., & Kanwisher, N. (1998). A cortical representation of the local visual environment. *Nature*, **392**(6676), 598–601.
- Gerbino, W., & Salmaso, D. (1985). Un'analisi processuale del completamento amodale. *Giornale Italiano di Psicologia*, **12**, 97–121.
- Gerbino, W., & Salmaso, D. (1987). The effect of amodal completion on visual matching. *Acta Psychologica*, **65**, 25–46.
- Grimson, W. E. L. (1981). *From Images to Surfaces. A Computational Study of the Human Early Visual System*. MIT Press: Cambridge, MA.
- Grossberg, S., & Mingolla, E. (1985). Neural dynamics of form perception: Boundary completion, illusory figures, and neon color spreading. *Psychological Review*, **92**, 171–211.
- Hochberg, J., & McAlister, E. (1953). A quantitative approach to figural "goodness." *Journal of Experimental Psychology*, **46**, 361–364.
- Huffman, D. A. (1971). Impossible objects as nonsense sentences. *Machine Intelligence*, **5**, 295–323.
- Kanizsa, G. (1955). Margini quasi-percettivi in campi con stimolazione omogenea. *Rivista di Psicologia*, **49**, 31–49.
- Kanizsa, G. (1976). Subjective contours. *Scientific American*. April.
- Kanizsa, G. (1979). *Organization in vision: Essays on Gestalt perception*. New York: Praeger. [see also Kanizsa, G. (1955). *Rivista di Psicologia*, **49**, 7].
- Kanizsa, G., & Gerbino, W. (1982). Amodal completion: Seeing or thinking? In J. Beck (Ed.), *Organization and representation in perception* (pp. 167–190). Hillsdale NJ: Erlbaum.

- 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 & Performance*. **24**(3),859–869.
- Koenderink, J. J. (1990). *Solid shape*. Cambridge, MA: MIT Press.
- Koenderink, J., & van Doorn, A. (1982). The shape of smooth objects and the way contours end. *Perception*, **11**,129–137.
- Koffka, K. (1935). *Principles of Gestalt psychology*. New York: Harcourt, Brace & Co.
- Kovacs, I. & Julesz, B. (1994). Perceptual sensitivity maps within globally defined visual shapes. *Nature*, **370**,644–646.
- Lowe, D. G. (1987). Three-dimensional object recognition from single two-dimensional images. *Artificial Intelligence*, **31**,355–395.
- Malik, J. (1987). Interpreting line drawings of objects. *International Journal of Computer Vision*, **1**,73–107.
- Michotte, A., Thinés, G., & Crabbé, G. (1964). Les complements amodaux des structures perceptives. *Studia Psychologica*. Louvain: Publications Universitaires. In G. Thinés, A. Costall, & G. Butterworth (Eds.), *Michotte's experimental phenomenology of perception*, Hillsdale, NJ: Erlbaum.
- Nakayama, K., & Shimojo, S. (1992). Experiencing and perceiving visual surfaces. *Science*, **257**,1357–1363.
- Nakayama, K., Shimojo, S., & He, Z.-J. (1995). Visual surface representation: A critical link between lower-level and higher-level vision. In Stephen M. Kosslyn & Daniel N. Osherson (Eds.), *Visual cognition: An invitation to cognitive science, Vol. 2. An invitation to cognitive science*, Cambridge, MA: MIT Press. (2nd ed., pp. 1–70).
- Nakayama, K., Shimojo, S., & Silverman (1989). Stereoscopic depth: Its relation to image fragmentation, grouping, and the recognition of occluded objects. *Perception*, **18**,55–68.
- Rensink R. A., & Enns J. T. (1995). Preemption effects in visual search: Evidence for low-level grouping. *Psychological Review*, **102**(1),101–130.
- Rensink R. A., & Enns J. T. (1998). Early completion of occluded objects. *Vision Research*, **38**, 2489–2505.
- Richards, W. A., Koenderink, J. J., & Hoffman, D. D. (1987). Inferring three-dimensional shapes from two-dimensional silhouettes. *Journal of the Optical Society of America A*, **4**(7),1168–1175.
- Rock, I. (1984). *The logic of perception*. Cambridge MA: MIT Press.
- Rubin, E. (1915/1958). Visuell wahrgenommene figuren. Copenhagen: Gyldenalske Boghandel, 1915. [reprinted as “Figure and ground” in D.C. Beardslee (Ed.), *Readings in perception* (pp. 194–203). Princeton, NJ: Van Nostrand, 1958].
- Sekuler, A. (1994). Local and global minima in visual completion: Effects of symmetry and orientation. *Perception*. **23**(5),529–545.
- Sekuler, A. B., and Palmer, S. E. (1992). Perception of partly occluded objects: A microgenetic analysis. *Journal of Experimental Psychology: General*, **121**,95–111.
- Sekuler, A. B., Palmer, S. E., & Flynn, C. (1994). Local and global processes in visual completion. *Psychological Science*. **5**(5),260–267.
- Sutherland, N. S. (1968). Outline of a theory of visual pattern recognition in animal and man. *Proceedings of the Royal Society, London, Series B*, **171**,297–317.

- Takeichi, H. (1995). The effect of curvature on visual interpolation. *Perception*, **24**,1011–1020.
- Takeichi, H., Nakazawa, H., Murakami, I., & Shimojo, S. (1995). The theory of the curvature-constraint line for amodal completion. *Perception*, **24**,373–389.
- Tse, P. U. (1998). Illusory volumes from conformation. *Perception*, **27**,8,977–994.
- Tse, P. U. (1999). Complete mergeability and amodal completion. *Acta Psychologica*. submitted.
- Tse, P. U., & Albert, M. K. (1998). Amodal completion in the absence of image tangent discontinuities. *Perception*, **27**,455–464.
- Ullman, S. (1990). 3-dimensional object recognition. *Cold Spring Harbor Symp. Quant. Biol.* **55**,889–898.
- Ullman, S. (1996). *High-level vision*. Cambridge, MA: MIT Press.
- van Lier, R. J, van der Helm, P. A., and Leeuwenberg, E. L. J. (1994). Integrating global and local aspects of visual occlusion. *Perception*. **23**(8),883–903.
- van Lier, R. J, van der Helm, P. A., and Leeuwenberg, E. L. J. (1995). Competing global and local completions in visual occlusion. *Journal of Experimental Psychology: Human Perception and Performance*. **21**(3),571–583.
- Waltz, D. (1975). Understanding line drawings of scenes with shadows. In P. H. Winston (Ed.), *The psychology of computer vision*. New York: McGraw-Hill.
- Wertheimer, M. (1923). Untersuchungen zur Lehre von der Gestalt. *Psychologische Forschung*, **4**,301–350.
- Wertheimer, M. (1938). *Laws of organization in perceptual forms*. London: Harcourt, Brace, and Jovanovich.
- Witkin, A., & Tanenbaum, J. (1983). On the role of structure in vision, chapter in: In Beck, Hope, and Rosenfeld (Eds.), *Human and machine vision*, Academic Press, New York. Pp. 481–543.
- Wouterlood, D., & Boselie, F. (1992). A good-continuation model of some occlusion phenomena. *Psychological Research/Psychologische Forschung*, **54**, 267–277.
- Yin, C., Kellman, P. J., & Shipley, T. F. (1997). Surface completion complements boundary interpolation in the visual integration of partly occluded objects. *Perception*, **26**(11), 1459–1479.

Received June 4, 1997