ELSEVIER

Acta Psychologica 102 (1999) 165–201

acta psychologica

Complete mergeability and amodal completion

Peter Ulric Tse [1]

Max Planck Institut fuer Biologische Kybernetik, 38 Spemannstrasse, 72076 Tuebingen, Germany

Received 29 June 1998; received in revised form 12 March 1999; accepted 13 March 1999

Abstract

When image fragments are taken to correspond to the visible portions of a single occluded object, the object is said to 'amodally complete' behind the occluder. Kellman and Shipley (Kellman, P. J., & Shipley, T. F. (1991). A theory of visual interpolation in objective perception. *Cognitive Psychology*, *23*, 144–221) argued that when the virtual contour extensions of such image fragments subtend an obtuse or right angle, the contours are 'relatable' and therefore complete. However, edge and surface relatability are neither necessary nor sufficient for completion to be perceived (Tse, P. U. (1999) Volume completion. *Cognitive Psychology*). Evidence is offered that completion is not driven directly by image cues such as contour relatability, but is driven, rather, by intermediate representations, such as volumes that are inferred from global image cue relationships. Evidence suggests that several factors, none of which is necessary for amodal completion to occur, contribute to the perceived strength of amodal completion, including similarity of pattern or substance, proximity, and good volume continuation or complete mergeability. Two partially occluded volumes are completely mergeable when they can be extended into occluded space along the trajectory defined by their visible surfaces such that they merge entirely with each other. Mergeability is not measurable in the image because it describes an inferred relationship among volumes that must themselves be inferred from the image. © 1999 Elsevier Science B.V. All rights reserved.

PsycINFO classification: 2323

Keywords: Vision; Form and shape perception; Perceptual closure; Visual perception

[1] Fax: 49-7071-601652; e-mail: tse@mpivision.kyb.tuebingen.mpg.de

PII: S 0 0 0 1 - 6 9 1 8 (9 9) 0 0 0 2 7 - X

1. Introduction

Several authors (e.g., Kanizsa, 1979, 1955; Michotte, Thinès & Crabbé, 1964/1991) have offered demonstrations that the visual system links image fragments into larger regions in a bottom-up or stimulus-driven manner (e.g., Fig. 1(a)). When those fragments are taken to be the visible portions of an occluded object, the object is said to 'amodally complete' behind the occluder. There has been a great deal of debate regarding the cues that the visual system uses to determine whether two partially occluded objects amodally complete into a single object behind an occluder. Broadly speaking, there have been three types of theories to account for amodal completion: (1) those that argue that completion occurs when certain conditions are satisfied among local image cues, such as contour orientations and junctions, (2) those that argue that completion occurs when certain conditions are satisfied among global image cues, such as symmetry, regularity, or simplicity of form or pattern and (3) those that argue that completion occurs when certain conditions are satisfied among internal representations, such as surfaces and volumes, which must be inferred from image cues.

More specifically, theories of type (1) have generally argued that completion occurs because of good continuation (Wertheimer, 1923) among image contours that terminate at an image tangent discontinuity such as a T-junction (e.g. Kellman & Shipley, 1991; Takeichi, Nakazawa, Murakami & Shimojo, 1995; Wouterlood & Boselie, 1992). Type (1) theories are essentially local in nature because the initiating conditions for contour interpolation are local tangent discontinuities. Initiated contour interpolations can link up at a distance with other interpolated contours if they meet the condition of good contour continuation. Theories of type (2) have challenged the notion that completion takes place because of local cues, such as contour orientations and discontinuities, and have emphasized instead the importance of global regularities such as symmetry in the patterns of completing image regions (e.g., Buffart, Leeuwenberg & Restle, 1981; Moravec & Beck, 1986; Sekuler, 1994; Sekuler, Palmer & Flynn, 1994; Van Lier, van der Helm & Leeuwenberg, 1994; Van Lier, van der Helm & Leeuwenberg, 1995). Most type (2) theories emphasize, global regularities in the inputs to completion processes, whereas others (Attneave &

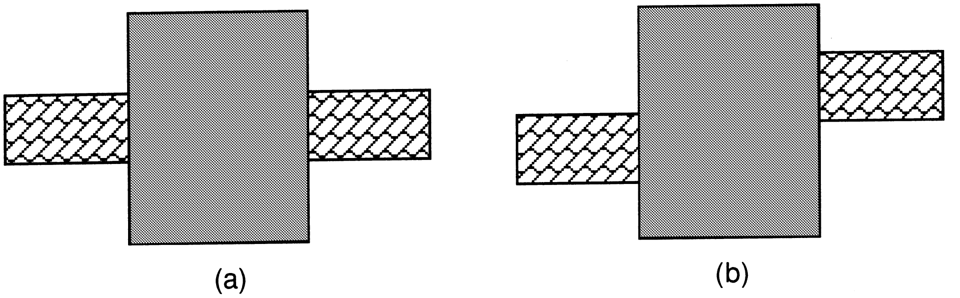

Fig. 1. Contours are relatable in (a) but not in (b). Amodal completion takes place in (a) but not in (b).

Arnoult, 1956; Buffart et al., 1981; Hochberg & Brooks, 1960) emphasize the importance of regularity and simplicity in the outputs of completion processes. Traditional theories of type (3) have emphasized that amodal completion occurs because of surface completion on a common depth plane (e.g. Nakayama & Shimojo, 1992; Nakayama, Shimojo & He, 1995; Nakayama, Shimojo & Silverman, 1989). A more recent type (3) approach has emphasized the importance of volume completion (Tse, 1998, 1999; Tse & Albert, 1998; Van Lier, 1999; Van Lier & Wagemans, 1999).

These three positions are not mutually exclusive. For example, according to the account of completion developed here, the visual system uses both (1) local and (2) global image cues to generate representations of (3) volumes that may occlude one another under constraint of (2) learned or innate priors that minimize output complexity. This paper describes empirical evidence for a new type (3) account of completion (Tse, 1998, 1999) that subsumes the flat surface approach as a degenerate case and emphasizes the importance of surface edges, curved surfaces, surface self-occlusions, and volumes in completion processes. According to this view, while the contour relationships emphasized by type (1) theories and global pattern regularities emphasized by type (2) theories play an important role in amodal completion, they are not direct cues to amodal completion. Rather, they and other image cues are used to infer edge, surface, and volume relationships in the world, and completion occurs at a level where these relationships are analyzed.

2. Contour and surface relatability

Kellman and Shipley (1991) argued that the image projection of partially occluded regions will amodally complete when their visible contours in the image are 'relatable'. Contours are relatable when they can be virtually extended so that their intersection subtends a 90° or larger angle in the image. Contours are relatable in Fig. 1(a), but not in Fig. 1(b). Thus, as predicted by Kellman and Shipley's theory of completion, the two partially occluded bars appear to be a single long bar in Fig. 1(a) but appear to be two separate bars in Fig. 1(b).

The proposal that completion is based upon contour relatability held broad appeal because one could determine whether contours were relatable just by comparing contour orientations in the image. A comparison of image contour orientations is the sort of algorithm that could be implemented in a reasonably straightforward manner using local operators. This suggested that one of the fundamental sources of image ambiguity, occlusion, might be overcome in a local and stimulus-driven manner, without recourse to complex global representations or higher-level knowledge. If image ambiguity due to occlusion could be overcome in a local, and stimulus-driven manner, perhaps other types of image ambiguity could also be overcome in this way. It might then be possible to maintain the reductionistic program of modeling vision as an entirely stimulus-driven succession of information processing stages over local cues.

It would indeed be computationally simplest if there were some locally measurable invariant cue in the image that was always present when the image depicted an occlusion scene and was never present otherwise. The visual system could then search for this occlusion cue and know with certainty that an object was occluded. Since Helmholtz (1962/1867) many researchers have maintained that image tangent discontinuities such as T-junctions approximate an invariant cue to occlusion, because these are always present in images depicting occlusion scenes (Clowes, 1971; Huffman, 1971; Kellman & Shipley, 1991; Lowe, 1985; Malik, 1987; Nakayama et al., 1989; Waltz, 1975). While it is true that T-junctions are generically present in the image when one surface occludes another surface separated in depth, T-junctions are not necessarily present in cases of perceived occlusion where one surface conforms to another (Tse & Albert, 1998). Similarly, it was suggested that contour relatability might also be an image cue generically present in all images projected from occlusion scenes (Kellman & Shipley, 1991). But this is also not true. Although contour relatability may be a useful cue to completion, it is neither necessary nor sufficient for a percept of amodal completion to take place (Tse, 1999; compare also Boselie & Wouterlood, 1992). For example, consider Figs. 2(a) and 2(b) reproduced from Tse (1999). In Fig. 2(a) completion occurs in the absence of relatable contours, and in Fig. 2(b), completion fails in the presence of relatable contours. Since the two most likely candidates for a necessary image cue to occlusion are T-junctions and relatable contours, and neither is a necessary or sufficient cue, it seems likely that no single image cue is necessarily present in the image projection of an occlusion scene. Likewise, there does not appear to be any image cue that is sufficient to infer a relationship of occlusion in the world. Thus a purely local image cue-driven account of completion is inadequate. Instead, completion should be understood to be a type of

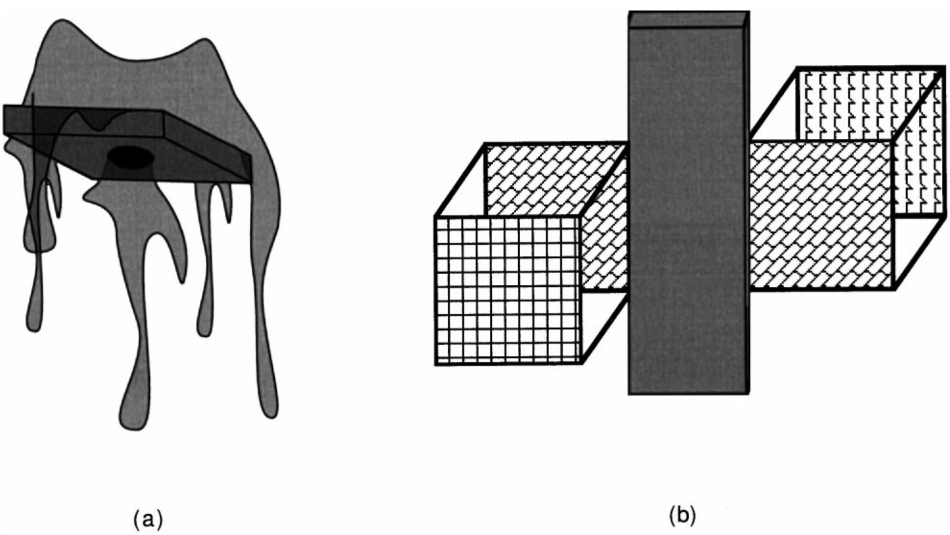

Fig. 2. In (a) amodal completion takes place in the absence of relatable contours or surfaces. In (b) amodal completion fails in the presence of both relatable contours and surfaces.

perceptual construction. Image cues are used to construct intermediate representations, and completion takes place over these representations.

Perhaps completion could be based on constructing and relating surfaces without having to go to a volume level of representation. However, even 'surface relatability' (among surface extensions into 3D occluded space) is neither necessary nor sufficient for amodal completion to be perceived (Tse, 1999). For example, in Fig. 2(a) there is amodal completion in the absence of both relatable edges or surfaces. That is, there are no partially occluded contours or regions (in the image) or edges and surfaces (in the world) that can be extended along their respective visible trajectories (in the image or into occluded 3D space, respectively) such that they will link up with other partially occluded contours, regions, edges or surfaces. Conversely, in Fig. 2(b), both relatable edges and relatable surfaces can be inferred, yet amodal completion fails, apparently because a surface cannot belong to two volumes at the same time.

These counter examples to 'traditional' contour- and surface-based theories of completion raise the fundamental question motivating the present investigation. If contour and region relatability in the image, as well as edge and surface relatability in the world, are not necessary and not sufficient for completion to take place, might there be some other more general principle that will allow us to predict when the visual system will amodally complete partially occluded volumes into whole volumes and when it will not? The proposed answer to this question for amodally completing surfaces and volumes is 'complete mergeability'.

3. Complete mergeability

An animal whose visual system completed occluded objects solely on the basis of image contour relatability might make fatal mistakes. Contours can be relatable in the image yet project from edges or surfaces that are not relatable in the world. Such a visual system might assume the existence of an occluded object when none was present, as in Fig. 2(b), or fail to complete a partially occluded object that really was a unified whole, as in Fig. 2(a), where there are no relatable contours. Similarly, an animal whose visual system only represented *visible* surfaces might make fatal mistakes because animals must move about and manipulate objects that self-occlude as well as occlude. For example, a monkey does not only swing from the visible surfaces of a branch, but must grasp according to the presumed form of the whole branch, including the occluded far side of the branch. The monkey must therefore have some internal representation of volumetric extent before it grabs the branch. Since an occlusion scene is generally comprised of 3D objects, a visual system that determined occlusion (including self-occlusion) relationships at a level where the spatial properties of objects are represented would not make such mistakes. In particular, volumetric extent and surface layout information for partially occluded objects is necessary to determine the extent and surface layout of the whole occluded object. For ecological reasons, it would behoove the visual system to use image cues to generate internal representations of volumes in order to determine the occlusion relationships that may exist among the volumes of the world.

In Tse (1999) it was argued that amodal completion takes place when partially occluded volumes are completely mergeable. 'Volume', by definition, means a bounded and uniformly connected enclosure such as a surface plus the inside space that it encloses. 'Complete mergeability' means that the virtual volume extension of a partially occluded volume can completely merge behind the occluder with the virtual volume extensions of volumes partially occluded by the same occluder. More specifically, 'complete mergeability' means that the unbounded surface and inside of a volume (or the interior of a flat surface in the degenerate case) can be extended into occluded space along the trajectory defined by its visible surfaces as they head under the occluder such that this unbounded volume extension merges entirely with other unbounded volume extensions. This can be thought of as 'good volume continuation', as opposed to good contour continuation. We shall see that several factors contribute to the strength of mergeability, including the proximity of partially occluded volumes, and the relative orientation of virtual volume extensions. How volume extensions behind an occluder merge will be explored in greater detail after several examples of amodal completion over volumes have been considered.

It must be stressed at the outset that complete mergeability, in contrast to contour relatability, is not measurable in the image and is not an image cue. This is because there are no volumes in the image. Volumes and their surfaces are inferred entities that must be constructed based on image cues. Complete mergeability is itself an inferred relationship among inferred entities derived from the image on the basis of various image cues to volume.

Fig. 3 depicts some examples of complete and incomplete mergeability for the degenerate case of surfaces that do not enclose space because they are flat. Since the principle of complete mergeability applies to both this degenerate domain and the more general domain of volumes, flat surfaces are depicted here because they are the traditional type of example used when discussing amodal completion and may be more familiar to the reader. Some corresponding demonstrations for the volume domain will be given later. The (a) side of this figure depicts the information available in the image. The (b) side depicts the type of information that might be inferred by the visual system given (a).

When a surface is partially occluded it can continue behind the occluder because it is unbounded and can link up behind the occluder with other unbounded surfaces (Nakayama et al., 1989, 1995; Nakayama & Shimojo, 1992). When there is only one unbounded surface as in Fig. 3(1a), that surface may extend, at least *in potentia*, under the occluder for some distance, as depicted in 3(1b). For the degenerate case of flat surfaces, the trajectory of the surface into occluded space is specified by the depth and trajectory of visible edge extensions into occluded space, providing overlap with traditional theories based on good contour continuation.

When a surface extension behind an occluder meets another surface extension such that their edges exactly coincide, as in Fig. 3(2b) or Fig. 3(6b) the probability that partially occluded surfaces will be perceived to complete is high because both extensions completely merge with the other. When a surface extension behind an occluder meets another surface extension such that some part of it does not merge with this other extension, it is 'partially unmergeable' with this other extension.

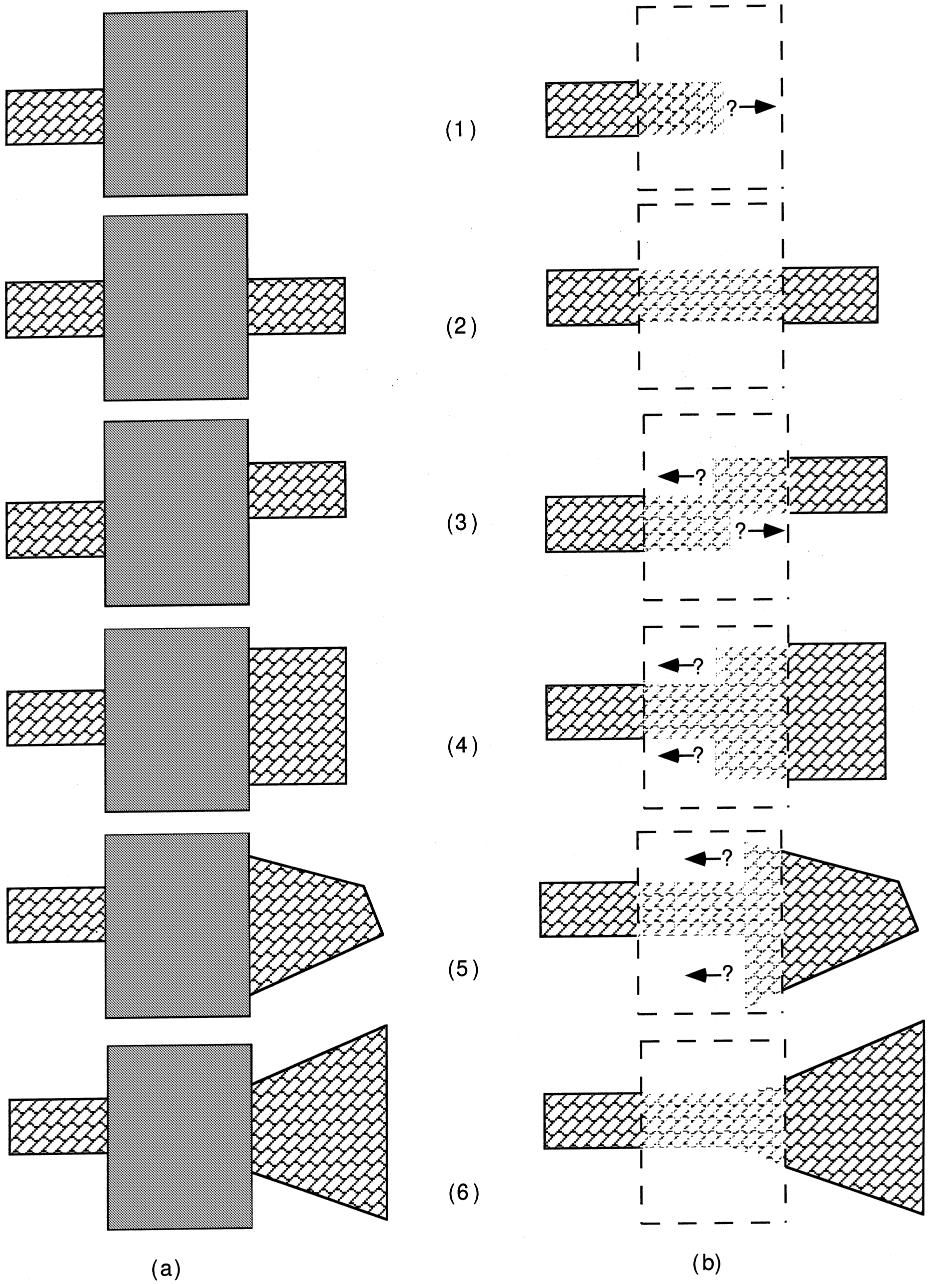

Fig. 3. Column (a) shows examples of occlusion, and column (b) shows how the partially occluded visible surfaces might extend into occluded space. In (1) the extension has nothing to link up with. Its extent is unspecifiable and may attenuate with distance. The traditional case of an amodally completing bar is shown in (2). Both contours are relatable and both bars are completely mergeable with one another in (2) and (6). In (3) complete mergeability fails both ways. In (4) and (5) complete mergeability is satisfied for the small bar, but is not satisfied for the large bar.

When an extension is partially unmergeable the extent of the unmerged portion behind the occluder is not determined. This is indicated in Fig. 3 with arrows and question marks. If two surface extensions meet such that both are partially unmergeable with the other, as in Fig. 3(3b), completion will generally fail, although this is a probabilistic claim, and the probability that completion will fail increases to the extent that complete mergeability is not satisfied. If only one extension is partially unmergeable, but the other is completely mergeable, as in Fig. 3(4b) and Fig. 3(5b), then completion will generally be perceived, albeit perhaps less compellingly than in cases where complete mergeability is satisfied both ways, as in Fig. 3(2a).

Amodal completion is not an all or nothing affair. For example, completion in Fig. 3(2a) may appear stronger than completion in Fig. 3(4a). If extensions behind an occluder attenuate with distance, it would seem reasonable to predict that the perceived strength of amodal completion will decrease with increases in the width of the intervening occluder. It would also seem reasonable to predict that the strength of perceived amodal completion will depend on the degree to which complete mergeability is satisfied. One might also wonder whether the length of the visible portions of a partially occluded surface has some influence on the degree to which it can extend into occluded space. These basic and preliminary issues were tested in Experiments 1 and 2 using stimuli from the traditional domain of flat surfaces.

4. General method: Experiments 1, 2, 4 and 5

Experiments 1, 2, 4 and 5 in this paper are actually the interleaved conditions of a single two alternative forced choice experiment run on 20 naive observers who had to respond whether two non-adjacent image regions appeared to amodally complete behind an occluder. Each observer was told: "Please decide whether there appears to be a single object behind the gray occluder or two separate objects behind the occluder. You may be able to see some stimuli either way. In this case, please respond according to which of the two percepts seems stronger to you. Indicate your percept by writing the first letter of your first name on the back of those sheets where the partially occluded regions appear to link up as a single object. Make no mark on the back of the sheet when the figure appears to be comprised of two separate occluded objects. Make your perceptual decision before turning the page, since other observers' marks may or may not be on the back of the page, and we do not want you to be biased by other observers' responses". Results were tallied by adding up the marks on the back of each page after all 20 observers had run.

Observers sat in a chair at a table and viewed the stimuli at rest on the table at a viewing distance of approximately 57 cm. Each trial consisted of a single gray scale image printed in the center of a standard 8.5 in. $\times$ 11 in. piece of white paper. The 178 trials for Experiments 1, 2, 4 and 5 were randomly interleaved. They were randomized by the author by blind shuffling the stimuli for 5 min before presenting the entire set of stimuli to each new observer. The author shuffled the stimuli by turning the stack of 178 sheets so that the side with the stimuli lay face-down. He then divided the stack into 20 piles of approximately equal numbers of sheets and then recombined these small stacks into a single large stack by picking them up in arbitrary order. This shuffling procedure was performed at least twice for each observer. The observer ran through 178 trials on 178 pieces of paper in a single sitting of approximately 30 min by flipping through the unstapled pages at his or her own pace. The observer could view each trial for as long as necessary though most observers took no more than a few seconds per trial to make a perceptual decision about amodal completion. The data will be described below separately as if they were separate experiments, since, at least in concept, they were. This 'interleaved' design was used so that new types of stimuli considered here would be judged in the context of traditional examples of amodal completion. The results of this experiment for the traditional domain will serve as a benchmark against which we can judge the strength of amodal completion in later experiments involving partially occluded volumes.

5. Experiment 1: testing contour relatability

This experiment tested the traditional claim that contour relatability is an important cue to amodal completion and also tested the possibility that the length of partially occluded visible surfaces contributes to the likelihood of amodal completion. Length was considered here to determine whether global factors play a role in this traditional domain, since global factors seem to play an important role in the volume domain, as suggested by Figs. 2(a) and 2(b).

5.1. Methods

Trials were randomly interleaved with those of Experiments 2, 4 and 5 as described above. A standard rectangular occluder (2.50×14.00 cm^2) and partially occluded bars of four lengths (1.00×0.62, 1.00×1.25, 1.00×2.50 and 1.00×5.00 cm^2) were tested, as shown in Fig. 4. Four bar lengths were tested in order to test the possibility that the longer the 'relating' visible bars, the stronger the percept of completion might be. The position of the left bar was held constant across trials. The position of the right bar for each bar length was varied around coalignment with the left bar, as shown in Fig. 1. Positions tested for the right-hand bar were 0.15, 0.30, 0.45, 0.60, 0.75, 0.90, 1.05, 1.35, 1.95 cm away from principal (central) axis coalignment in both the upward and downward directions. Including the coaligned position there were 19 positions tested per condition. The amount of displacement between bars shown for the $4x$ case was the same used in the $1x$, $2x$ and $8x$ bar length cases, respectively.

5.2. Results and discussion

When the left and right bars were coaligned, amodal completion was at or near 100% for all four bar lengths, as would be expected purely on the basis of contour or edge relatability. As the bars deviated from coalignment, the percentage of observers

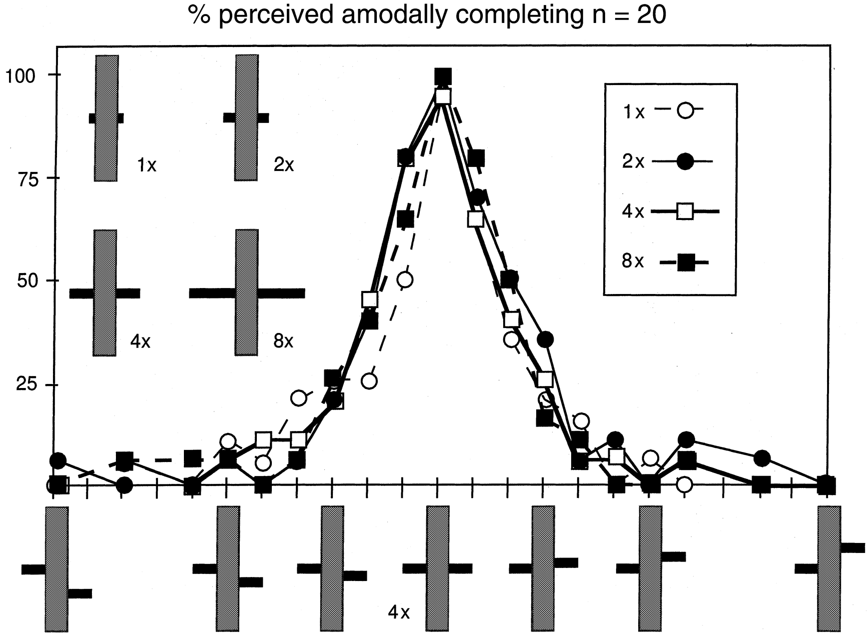

Fig. 4. The y-axis indicates the percentage of subjects who saw the configurations shown on the x-axis as amodally completing. Right-hand extensions were varied around coalignment (center case). Four bar lengths were tested. The displacement around coalignment was the same for each extension length.

who reported seeing the two bars amodally complete fell off rapidly to approximately zero. The data for the four bar length conditions was tested to see whether the four sets of bar length data shown in Fig. 4 were statistically drawn from a common population or a different population of scores. Two χ^2 tests were performed on the data. The first compared observed versus expected percentages for the central 13 displacements along the x-axis. The data for the three displacement values at either end of the x-axis were discarded because all curves in these outlying regions seemed to lie in the neighborhood of 0%, and any differences between conditions would be most detectable in the region of small displacements away from coalignment. The result of the χ^2 test was equal to 21.5, which at 36 degrees of freedom (d.f.), was far from significant. A second χ^2 test was performed by examining only data for the central nine displacement values, since these were judged to be the region where a significant difference, if there was one, might be found. The result of this χ^2 test was equal to 10.3, which at 24 d.f., was also far from significant. The null hypothesis that all data points were drawn from a common population of scores could therefore not be rejected. Thus the length of the bars did not play a significant role in determining whether observers would report the bars to amodally complete in these configurations. It would seem that only the alignment of partially occluded edges in the immediate neighborhood of the occluder were used by amodal completion processes.

6. Experiment 2: measuring the effect of occluder width

This experiment examined whether the width of the occluder plays a role in amodal completion as predicted by the idea that surface extensions into occluded space attenuate with distance, as depicted in Fig. 3(b). If the strength of extensions attenuates with distance, we would expect the likelihood of amodal completion to decrease with increasing occluder width, holding the amount of bar displacement from coalignment constant.

6.1. Methods

Trials were randomly interleaved with those of Experiments 1, 4 and 5, as described above. Three rectangular occluders were used for each bar displacement value. The widths were 1.25×14.0, 2.50×14.0 and 5.00×14.0 cm^2. Only a single bar length of 1.00×0.62 cm^2 was tested, as shown in Fig. 5. The $2x$ width case here was the same as the $1x$ bar length case of Experiment 1, and the same data is shown in Fig. 5 for this case. The position of the left bar was held constant across trials. The position of the right bar for each bar length was varied around coalignment with the left bar, as shown in Fig. 5. Positions tested for the right-hand bar were 0.15, 0.30, 0.45, 0.60, 0.75, 0.90, 1.05, 1.35, 1.95 cm away from coalignment in both the upward

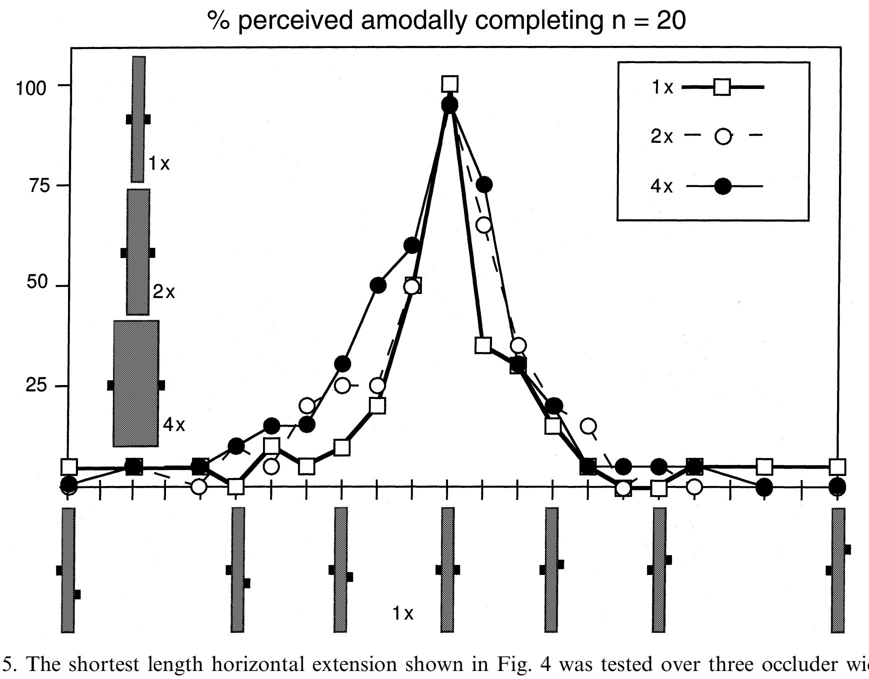

Fig. 5. The shortest length horizontal extension shown in Fig. 4 was tested over three occluder widths. Displacements of the right extension relative to the left one are the same for all occluder widths. The position of the left extension was held constant.

and downward directions. Including coalignment, there were thus 19 tested positions per condition. The same amount of displacement was used for all occluder width levels.

6.2. Results and discussion

Again, when the left and right bars were coaligned, amodal completion was at or near 100% for all three occluder lengths, as would be expected purely on the basis of contour relatability. As the bars deviated from coalignment, the percentage of observers who reported seeing the two bars amodally complete fell off rapidly to approximately zero. The data for the three occluder width conditions was tested to see whether the three sets of occluder width data shown in Fig. 5 were statistically drawn from a common population or a different population of scores. Two χ^2 tests were performed on the data as in Experiment 1. The first compared observed versus expected percentages for the central 13 displacements along the x-axis. The result of the χ^2 test was equal to 13.3, which at 24 d.f., did not reach significance. A second χ^2 test was performed by examining only data for the central nine displacement values. The result of this χ^2 test was equal to 8.5, which at 16 d.f., was also not significant. Thus occluder width does not appear to play a significant role in determining whether observers report the bars to amodally complete, at least for the configurations considered here.

In general the data from Experiments 1 and 2 support the traditional view that amodal completion is strongly dependent on contour relatability in the image or edge relatability in the world. Note, however, that contour relatability is confounded in this stimulus configuration with complete mergeability. That is, extending the interiors of the two partially occluded bars (rather than just their contours) along the trajectory defined by their visible contours, will lead to a complete merging of those interiors. In order to test whether complete mergeability makes a contribution to amodal completion that is independent of contour relatability, stimuli that satisfy complete mergeability but lack contour relatability were constructed for Experiments 3, 4 and 5.

7. Experiment 3: isolating complete mergeability

The null result of Experiment 2 is perhaps surprising because it seems unlikely that extension strength would continue unabated with distance. If the occluder were many times wider than the $4x$ width tested, perhaps the likelihood of amodal completion would decrease. This null result may be due to a variety of factors. First, the strength of an extension behind an occluder may only fall off beyond a certain threshold of distance and the occluder widths tested in Experiment 2 were insufficient to probe this. Second, the visual system may be highly sensitive to coaligned contours even over large visual angles. Perhaps if observers had to respond to stimuli that lacked coaligned contours their responses would better reflect the presumed attenuation of surface extensions behind an occluder. Lastly, the two alternative forced choice method used in Experiments 1 and 2 may not be a sufficiently sensitive experimental paradigm to probe the perceived strength of amodal completion. Perhaps observers set a low threshold for saying that they perceived amodal completion. A more sensitive experimental paradigm might therefore be a task where observers rate the perceived strength of amodal completion among stimuli that lack coaligned visible contours. This is the paradigm used in the present experiment.

While it is possible to 'will' to see amodal completion between the rectangle and triangle in Fig. 6, this appears to be a fairly low probability configuration, as will be demonstrated in this experiment. We can imagine the unbounded sides of the rectangle and triangle extending into occluded space at a common 'rate'. In order for the rectangle in Fig. 6 to merge with the triangle it would have to extend far enough to merge with the triangle 'before' the triangle closed at the point defined by the intersection of its two unbounded contour extensions. While there are no grounds to assert that the visual system propagates surface extensions over space and time in this manner, this might offer a reasonable rule of thumb for estimating the strength of perceived amodal completion. Configurations where one partially occluded surface can close up 'before' another unbounded surface extension can merge with it may be configurations where amodal completion is weak if it is perceived at all.

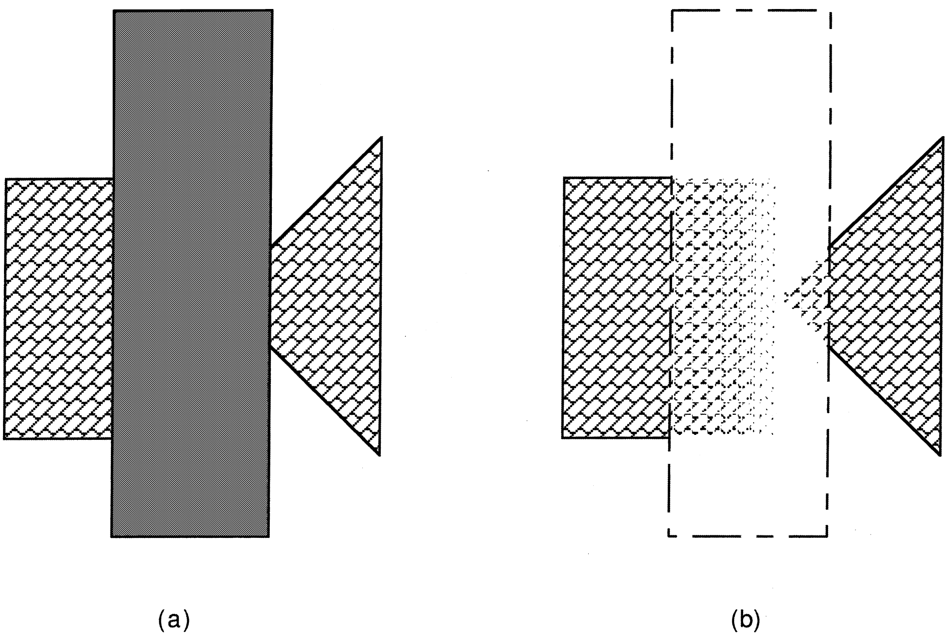

Fig. 6. Amodal completion between the partially occluded triangle and rectangle may appear weak here if it takes place at all.

In Fig. 7(a), amodal completion may appear to take place more compellingly than in Fig. 6(a). This may be because the rectangle can merge with the triangle of Fig. 7(a) 'before' it closes up, whereas this is not true for the triangle of Fig. 6(a), assuming a common rate of extension. If the two surfaces extended behind the occluder at the same rate they would meet at approximately the configuration depicted in Fig. 7(b). If they extended at different rates, such that the rectangle extended 'faster' than the triangle, the configuration might look something like Fig. 7(c). While this is not an impossible solution, it may be less probable than the solution shown in Fig. 7(b).

Fig. 7(a) might undergo amodal completion more compellingly than Fig. 6(a) for another reason as well. The deepness of the concavity where the triangle's presumed extension meets the presumed extension of the rectangle is greater in Fig. 6(a) than in Fig. 7(a). Several authors (e.g. Hoffman & Richards, 1984; Hoffman & Singh, 1997) have suggested that we divide objects into parts at points of deep concavity along the bounding contour, and that the likelihood of dividing objects into parts increases with the deepness of the concavity where a potential part meets the whole. Thus, at least two possible mechanisms may account for why Fig. 7(a) appears to undergo amodal completion more readily than Fig. 6(a). This experiment will show that both factors make significant contributions to the perceived strength of amodal completion.

7.1. Methods

Eight naive observers viewed the stimuli shown in Fig. 8. The stimuli were comparable in size to those used in previous experiments, and each of the 24 configurations used was printed in gray scale in the center of an 8.5 in. $\times$ 11 in. piece of paper as in Experiments 1, 2, 4 and 5. Three occluder widths were tested. The widths were 1.30×11.50, 2.60×11.50 and 3.90×11.50 cm^2. The size of the left-hand bar was held constant at 2.7×4 cm^2 and its position was held constant at the center of the left side of the occluder. The black surface on the right side of the occluder was drawn so that it might be the visible portion of an isosceles triangle. Moving from left to right in Fig. 8, the common angle of the visible portions of the isosceles triangle was 90°, 81°, 78°, 73°, 60°, 50°, 40° and 30°. The width of the visible portion of the virtual isosceles triangle was held constant at 4 cm. Stimuli were randomized by blind shuffling the papers for 1 min for each new observer. Each trial consisted of a single page. Observers viewed the stimuli at a table from a viewing distance of approximately 57 cm. Instead of performing a two alternative forced choice task as in the other experiments in this paper, observers responded by assigning an integer between 0 and 10 corresponding to the perceived strength of amodal completion for each trial. 10 was assigned if the triangle and rectangle definitely appeared to be a single object behind the occluder. 0 was assigned if they definitely appeared to be two different partially occluded objects. On each of the 24 pieces of paper was written a number between 1 and 24. Observers recorded their response next to the corresponding number on a separate sheet of paper that had the numbers 1 through 24 printed on it corresponding to the 24 trials.

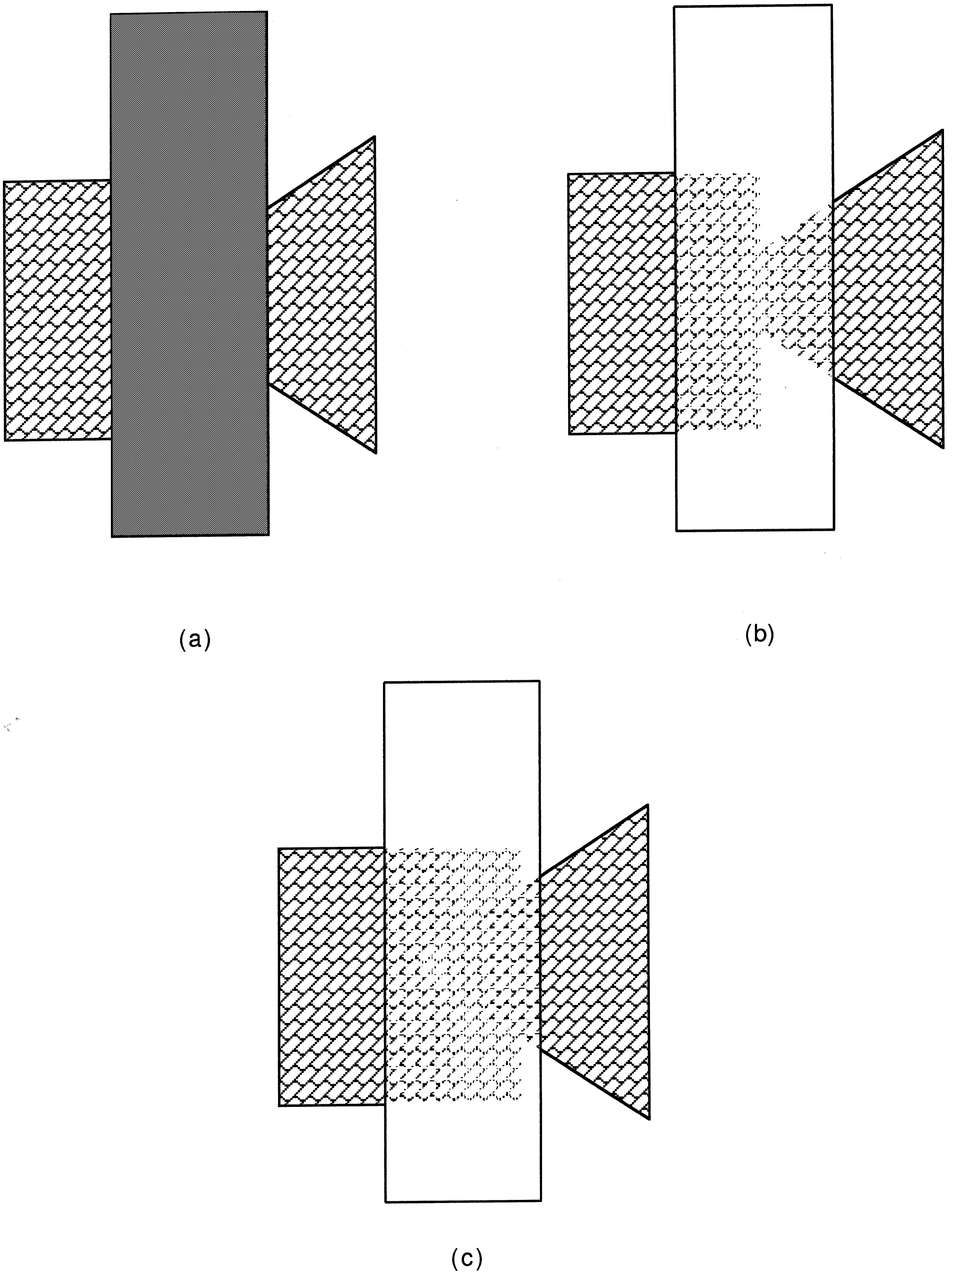

Fig. 7. Amodal completion in (a) may appear stronger than in Fig. 6 for reasons described in the text. If the surfaces extend into occluded space at the same 'rate' they might merge as in (b). If they extend at different rates they might merge as in (c).

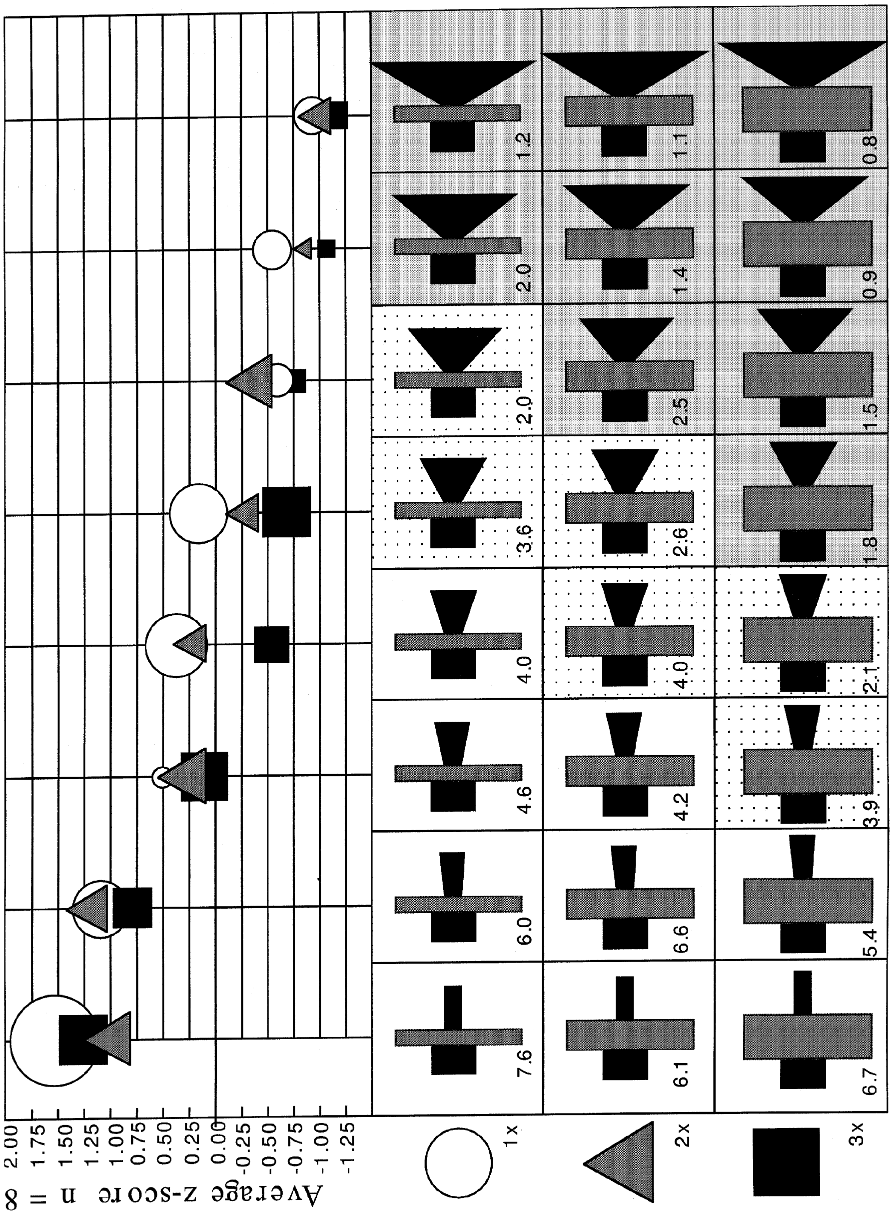

Fig. 8. All stimulus configurations used in the experiment are shown below the graph. The raw score obtained for each stimulus is shown to its lower left. The white circle, gray triangle and black square indicate stimuli in the top, middle and bottom row, respectively. Standard errors around the mean for Z-scores are indicated by the top and bottom points of the circle, triangle and square.

7.2. Results and discussion

Normalized scores are shown in the graph in Fig. 8. Z-scores are shown in the figure because observers utilized the 0 to 10 scale differently. The white circles, gray triangles and black squares correspond to the thin, medium and wide occluder cases respectively. The top and bottom of each white circle, gray triangle, or black square in the graph indicate the locations of one standard error around the mean.

An ANOVA (total d.f. = 191, error d.f. = 98) was performed on raw scores that considered up to two-way interactions among factors. Factors that absorbed variance were triangle shape (TS), observers (Ob) and width (W) of the occluder. Average raw scores are shown to the lower left of each case in Fig. 8. $F(7,98)$ for the Ob factor was 14.78 ($p < 0.0001$). $F(7,49)$ for the TS factor was 27.82 ($p < 0.0001$). Thus, as we move from left to right along each row, the likelihood that an observer will report the partially occluded black triangle and rectangle to amodally complete behind the occluder into a single object behind the occluder decreases significantly. This 'column effect' might be due to at least two factors. First, as we move from left to right, the sharpness of the deep concavity that the extension of the partially occluded triangle makes with the extension of the partially occluded rectangle increases. Second, as we move from left to right along each row, the extension of the partially occluded triangle closes into a triangle further towards the right-hand side of the occluder. We can test for the relative contributions of these two possible causes of the column effect by comparing the three rows.

For each column comprised of the three cases where the triangle has a given visible angle, the angle of the concavity made by the extension of the triangle with the extension of the rectangle remains constant. However, as we move from the top (white circle) row down to the bottom (black square) row, the extension of the partially occluded triangle closes further to the right because the occluder is greater in width. For example, in stimuli with a white background in Fig. 8, the virtual triangle extension would close beyond the left side of the occluder and therefore definitely merge with the rectangle. In the cases with a dotted background, the triangle extension would close behind the half of the occluder closer to the visible portion of the partially occluded rectangle. In the cases with a uniform gray background, the triangle extension would close behind the half of the occluder closer to the visible portion of the partially occluded triangle. (In fact, all stimuli were shown on a white background). If 'speed of closure' makes a significant contribution to the perception of amodal completion we would expect there to be a 'row effect', such that the wider the occluder the less likely amodal completion would be perceived for a given column. $F(2,14)$ for W was 8.48 ($p = 0.0039$).

Only one of the three two-way interactions among the factors reached significance. $F(49,98)$ for Ob×TS was 2.99 ($p < 0.0001$). Interactions that did not reach significance were as follows. $F(14,98)$ for Ob× W was 1.65 ($p = 0.078$), and $F(14,98)$ for W×TS was 1.40 ($p = 0.166$). Thus observers utilized the subjective scale significantly differently for the column effect, but not for the row effect.

This experiment shows that complete mergeability makes a contribution to the percept of amodal completion that is not reducible to an effect of contour relatability. Moreover, it shows, in contradiction to the null result of Experiment 2, that occluder width and/or ‘speed of closure’ makes a significant contribution to the perceived strength of amodal completion.

8. Experiment 4: complete mergeability in the absence of coalignment

Coaligned bars of the same width, such as the center cases in Experiments 1 and 2, have both completely relatable contours and are completely mergeable with one another. Coaligned identical bars thus do not provide a crucial test between the predictions made by the principles of complete mergeability, on the one hand, and image contour relatability, on the other. In Experiment 3, column 1 of Fig. 8 consisted of bars of two different widths. Amodal completion was perceived to a high degree despite the absence of relatable contours. Since this configuration satisfies the principle of complete mergeability in the absence of relatable contours, it seems that complete mergeability can give a more general account of amodal completion than can contour relatability. However, the three examples in the first column of Fig. 8 are nonetheless coaligned, because they share a common principal or central axis. If the principle of complete mergeability is correct, even when the principal axes are not coaligned, completion should take place as long as at least one of the bars is completely mergeable with the other. This experiment shows that this is indeed the case.

8.1. Methods

The stimuli used in this experiment were randomly interleaved with the stimuli used in all experiments except Experiment 3. The occluder was the same as that used in Experiment 1 (2.50×14.00 cm^2). The position of the left bar ($1.00°\times2.50°$) was held constant at the center of the left side of the occluder, while the right-hand bar (3.70×2.50 cm^2 or 1.85×2.50 cm^2) was varied in distance from the center of the right side of the occluder. Positions tested for the right-hand bar were 0.15, 0.30, 0.45, 0.60, 0.75, 0.90, 1.05, 1.35, 1.95 and 2.55 cm away from the center of the right side of the occluder.

8.2. Results and discussion

The data for the case where the right-hand bar is 3.70×2.50 cm^2 are shown in Fig. 9. Those for the case where the right-hand bar is 1.85×2.50 cm^2 are shown in Fig. 10. In both cases the left-hand bar can undergo complete mergeability with the right-hand bar as long as the upper (lower) edges of both bars are relatable to one another. These points are indicated with arrows in both Figs. 9 and 10. Note that all positions between these two extremes were seen to amodally complete by a majority of observers. Complete mergeability follows because the left-hand bar's interior can fully merge with that of the right-hand bar, even in the absence of contour or edge relatability. Once the left-hand bar's upper (lower) edge gets above (below) that of the right-hand bar's upper (lower) edge, extending the left-hand edge to the right would

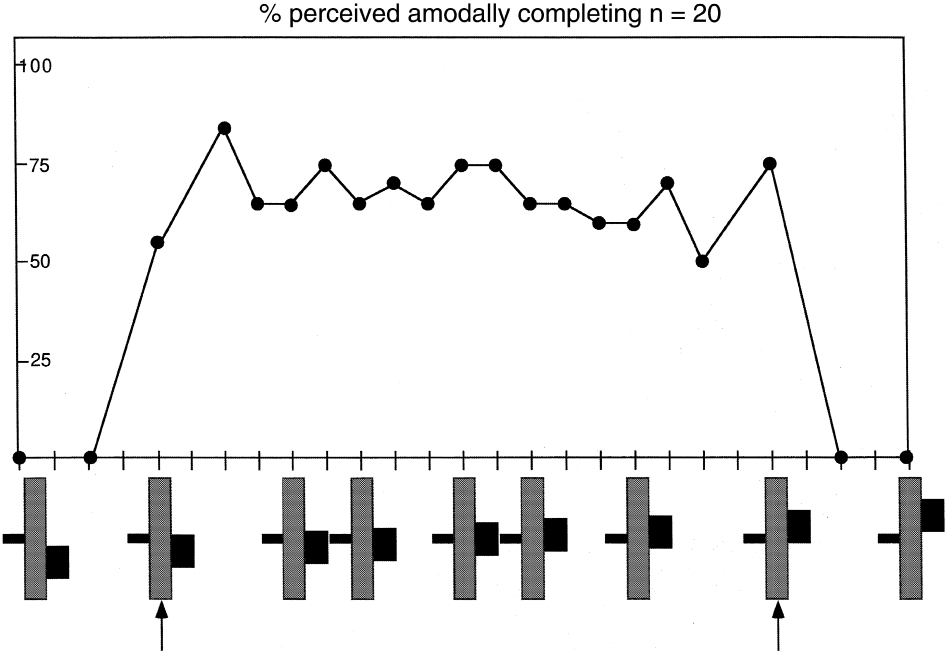

Fig. 9. Amodal completion in the absence of contour relatability takes place because of complete mergeability. The arrows indicate stimuli where there is both partial contour relatability and complete mergeability. All stimuli between the arrows are mergeable from left to right.

not lead to a complete merging of all the left-hand bar's interior with that of the right-hand bar's interior.

Note that complete mergeability only works here for the left-hand bar; Extending the right-hand bar to the left does not lead to a complete merging of its interior with that of the left-hand bar's interior. This does not eliminate the possibility of amodal completion since it only appears necessary that complete mergeability work in at least 'one direction'. Stimuli to the left of the left arrow or to the right of the right arrow failed complete mergeability both ways. Not surprisingly, the percentage of observers who reported seeing a single object behind an occluder in these configurations was close to zero. When both bars' edges are entirely aligned on top and bottom, complete mergeability works in both directions, and may account for the greater strength of perceived amodal completion in the traditional cases shown below the central x-axis position in Figs. 4 and 5. While the strength of amodal completion in Figs. 9 or 10 never reaches the 100% strength found in the case of perfectly aligned edges or contours shown in Fig. 4, complete mergeability appears to make a contribution to amodal completion that is not reducible to contour relatability. Indeed, contour relatability appears to fail to generate a percept of amodal completion when complete mergeability fails. For example, when the upper contour of one bar is relatable to the lower contour of another, as in the second $4x$ case from the left in Fig. 4, amodal completion is not perceived. Thus, amodal completion is

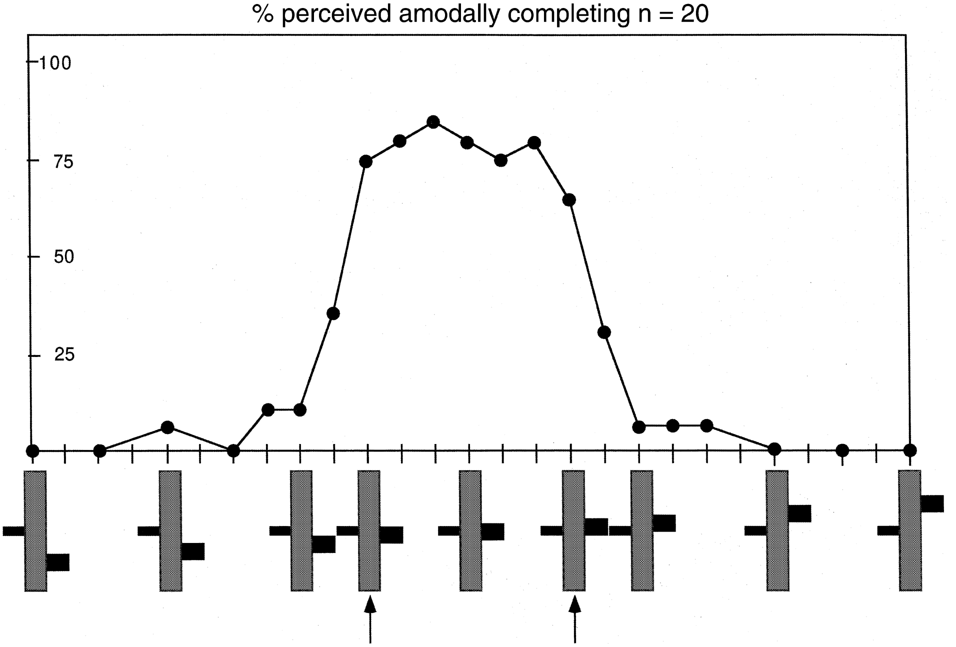

Fig. 10. The same experiment as shown in Fig. 9 was carried out with a narrower right-hand extension. All stimuli at and between the arrows are mergeable from left to right.

perceived when there is complete mergeability in the absence of contour relatability, but amodal completion is not perceived when there is contour relatability in the absence of complete mergeability. This suggests that surface or volume mergeability rather than contour (or edge) relatability is the more fundamental principle governing amodal completion over surfaces and volumes.

It could be that contour or edge relatability would enhance a percept of amodal completion given that the principle of complete mergeability is satisfied. Note, for example, that more observers reported a percept of amodal completion in the case in Fig. 9 indicated with the right-hand arrow. In this case the lower contours of both the wide and narrow partially occluded bars coalign. However, this did not hold true for the case indicated by the left-hand arrow. Moreover, there was no apparent contribution of contour relatability above and beyond that made by complete mergeability in the two corresponding cases indicated with arrows in Fig. 10. Overall, then, it does not seem that contour relatability enhances the strength of completion, although it may do so weakly.

9. Experiment 5: complete mergeability in the volume domain

The examples shown in figures up to this point have involved amodal completion between flat surfaces that appear to lie at the same depth. The examples described in

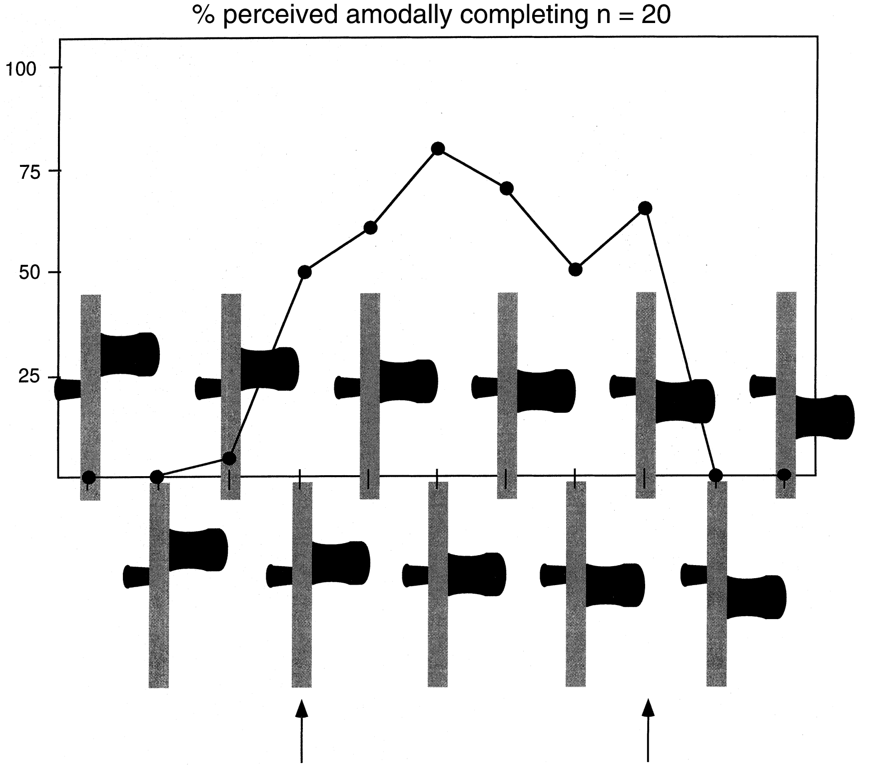

Fig. 11. The visible portions of the partially occluded surfaces appear to be volumetric pegs. All stimuli at and between the arrows are mergeable from left to right.

Experiment 5, shown in Figs. 11–13, are drawn to appear volumetric. That is, their surfaces appear to enclose space. Thus, in the volume domain complete mergeability can refer not only to the merging of the interior regions of surfaces, as in previous experiments, but also to the merging of the insides of a volume (Tse, 1999).

9.1. Methods

These stimuli were randomly interleaved with the stimuli for Experiments 1, 2 and 4. The dimensions of the occluder used for the data set shown in Fig. 11 was 1.25×13.20 cm^2 The width of the left-hand 'peg' was 1.60 cm at its widest, and the right-hand peg was 3.60 cm wide at its widest. The height of the left-hand peg was 1.4 cm and of the right-hand peg was 2.2 cm at its highest. The position of the left-hand peg remained fixed at the center of the left-hand side of the occluder for all trials. The position of the right-hand peg was displaced from the center of the right side of the occluder by 0.00, 0.35. 0.71, 1.17, 1.68 and 2.20 cm in the upward and downward directions for a total of 11 trials.

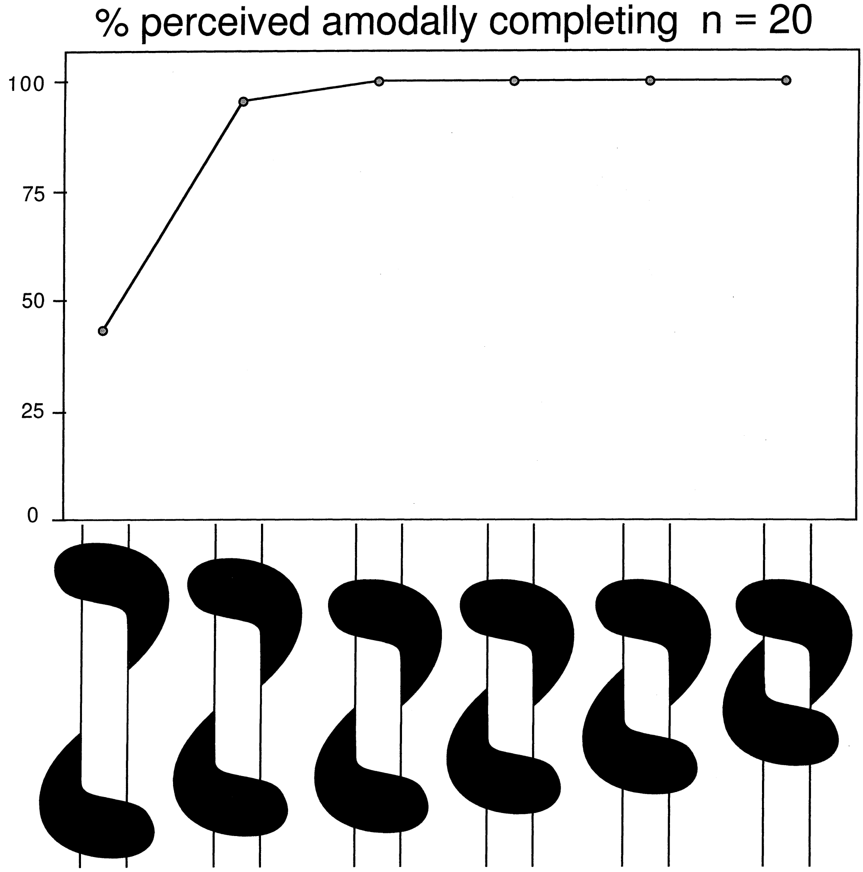

Fig. 12. The percentage of subjects who saw the two halves of the 'worm' complete for various configurations shown along the x-axis.

The six stimuli shown in Fig. 12 consisted of identical left-hand and right-hand portions of a partially occluded 'worm'. Both portions were moved up and down around the center of the 'pole' as shown in the figure.

The five stimuli shown in Fig. 13 consisted of a drawing of a cylinder whose front surface was 14 cm high and 2.4 cm wide. The position of the left-hand blob was fixed, while that of the right-hand blob was varied as shown in the figure.

9.2. Results and discussion

The case shown in Fig. 11 replicates the flat surface examples of Experiment 4, except that the amodally completing image regions are drawn as silhouettes that evoke a percept of cylindrical pegs for most observers. Amodal completion here seems to occur with roughly the same strength as in the flat surface examples of Figs.

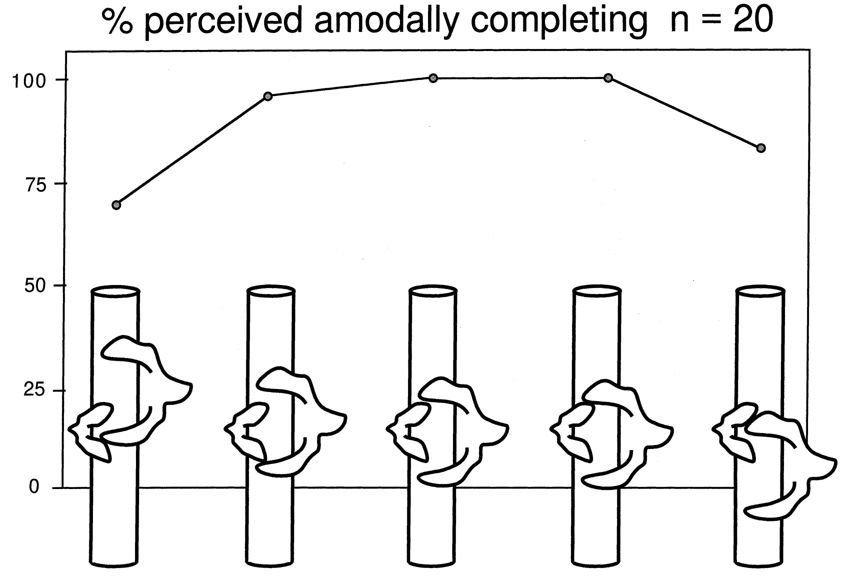

Fig. 13. The percentage of subjects who saw the two blobs complete for various configurations is shown along the x-axis.

9 and 10 even though the left-hand and right-hand pegs cannot have surfaces that lie at the same depth as each other if they are seen to be volumetric. If the pegs are the same size, then they must lie at radically different depths. If they are different sizes, and their principal axes lie at a common depth, then at least some portions of the larger peg are not completely mergeable with the smaller peg, assuming that the large peg extends behind the occluder along the trajectory specified by its visible curved surfaces. However, because of the nature of these stimuli it is possible that the black visible surfaces in Fig. 11 are still being treated as flat rather than curved surfaces. The following examples, however, force the visual system to treat the visible surfaces as curved surfaces.

The cases shown in Fig. 12 are further examples of amodal completion in the absence of relatable contours. Here observers generally reported seeing a 'worm' wrapped around a pole. The worm completes despite quite radical displacements of the two visible portions of the worm. The worm that completes is considerably thicker in the rightmost case than in the leftmost. In the second case from the left in Fig. 12 the contours of the two visible portions of the worm are actually relatable as contours. Indeed, for 95% of observers the two portions here were perceived to link up. But the visual system does not seem to link these two contours up per se, because the resulting single contour would violate complete mergeability. This follows because the interiors of the two extensions would lie on opposite sides of that linked contour and therefore fail to merge. Instead, the visual system appears to prefer to link the two visible portions together as a very thin volume. In this case the two relatable contours of the two visible portions interpolate a very thin tube rather than a single contour. Remarkably, even for the leftmost case, where the contours of the two visible portions of the worm are not relatable, nearly half of all observers reported seeing the two visible portions interpolate a very thin tube behind the occluder. For ambiguous cases, this may be because the visual system tends to assume connectedness as its default, as described in the next section.

In Fig. 13 it is not possible to treat the visible portions as flat surfaces because they appear to have self-occlusions, and flat surfaces cannot have self-occlusions. These occlude any contours that might have been relatable. This, then, is another example of amodal completion in the absence of contour relatability. The data in Fig. 13 reveal the remarkable flexibility of the amodal completion process. Amodal completion takes place in each case for most observers despite quite radical displacements of the partially occluded volumes relative to one another along the cylindrical occluder. Similar displacements in the flat surface domain would abolish the percept of amodal completion, as the data in Fig. 4 make clear. However, here, in the volume domain, the visual system seems quite liberal in assuming amodal completion. Indeed, even the extreme ends of this stimulus continuum are perceived to amodally complete with a higher proportion than the optimal stimulus configurations found in Figs. 9 or 10.

10. Defaults of completion

It may be that the visual system assumes amodal completion as its default, in the absence of image contour cues that expressly dispel that possibility (Tse, 1999). On this ‘connected unless proven disconnected’ view, image cues are cues of denial rather than cues of permission. Such a view would account for the finding here that the interpolated form of the entire volume is very different in each of the five cases shown in Fig. 13 even though the two amodally completing visible portions are the same in each case. In Fig. 13 there are no relatable contours or contours that are partially occluded by the occluder. Therefore, none of the visible contours can act as cues of permission for amodal completion. As mentioned before, it is possible that contours could also act as cues of permission if they actively enhanced the perceived strength of amodal completion rather than just failed to inhibit it. It was noticed in the discussion of Fig. 9 that where both complete mergeability and contour relatability are satisfied, contour relatability may have a weak enhancing effect on the perceived strength of amodal completion, although this was not observed in Fig. 10. Image contours and other cues may therefore serve a dual role, as cues of permission and cues of denial. The evidence presented in this paper, however, suggests that they serve primarily as cues of denial.

There can be positive evidence that two image regions correspond to connected objects in the world and there can be positive evidence that they correspond to disconnected objects. Positive evidence that two image regions (that abut a common region projected from their presumed common occluder) correspond to a single connected object may be any relevant commonality of image information or information inferred from image information. Conversely, positive evidence that two regions correspond to disconnected objects may be any relevant difference of image information or information inferred from image information. In an idealized analysis, there would be situations where the visual system has no positive evidence that two things are connected, and there would be other situations where the visual system has no positive evidence that they are disconnected. The visual system can make four possible assumptions in these circumstances. (1) When there is no positive evidence that they are disconnected, assume that they are connected by default. (2) When there is no positive evidence that they are disconnected, assume that they are

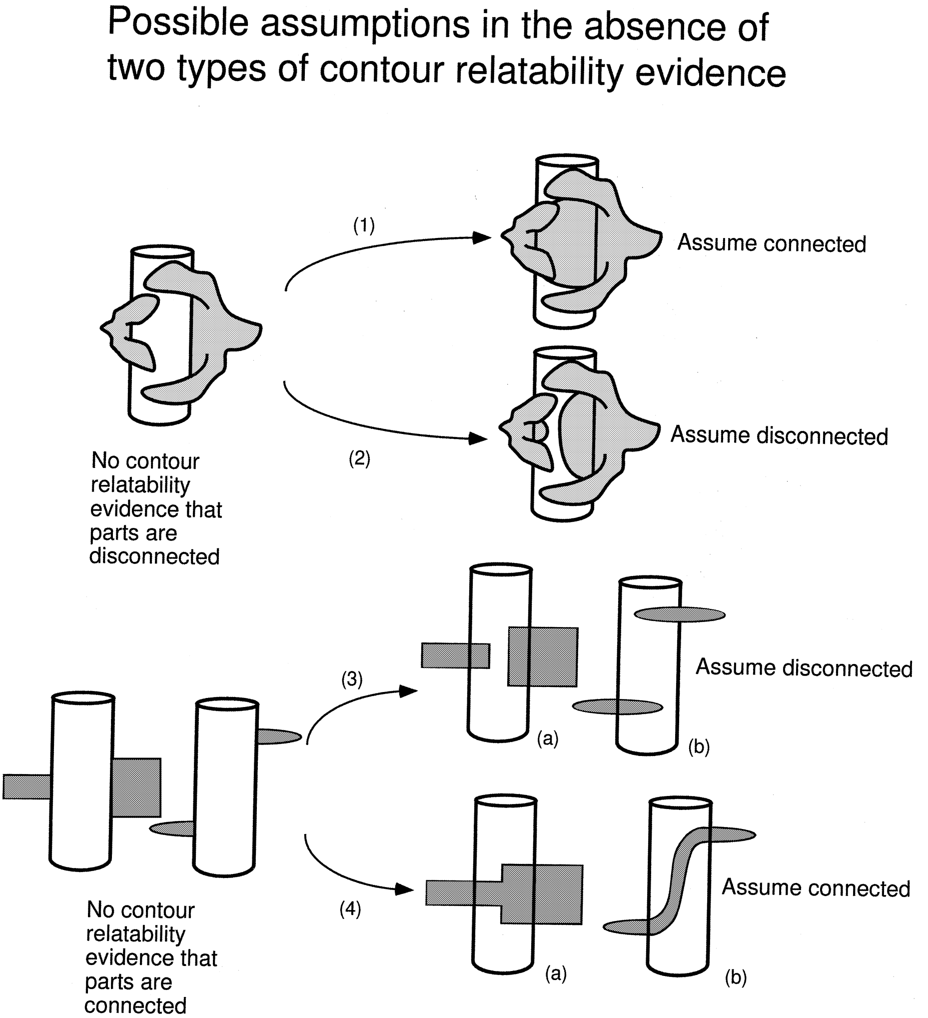

Fig. 14. Image cues can serve as cues of denial or permission for completion. Various possible assumptions that the visual system might make in the absence of cues of denial or cues of permission are schematized here and discussed in the text.

disconnected by default. (3) When there is no positive evidence that they are connected, assume that they are disconnected by default. Or finally, (4) when there is no positive evidence that they are connected, assume that they are connected by default.

An example where there would be no positive contour evidence that they are disconnected would arise when there are no visible relatable contours in an image cast from an occlusion situation where even if relatable edges were present in the world, they would not be visible. This is the situation in Fig. 13, and is shown again at the top of Fig. 14. A graphic depiction of assumptions (1) and (2) is drawn for this case. The data in Fig. 13 make clear that assumption (2) is not made by the visual system.

An example where, at least according to the contour relatability view, there would be no positive contour evidence that the partially occluded objects are connected would arise when there are no relatable contours in an image of an occlusion scene, but where all relevant edges are visible as they head under the occluder. A graphic depiction of assumptions (3) and (4) is drawn for this case in Fig. 14. Since the visual system completes as in (3b), obeying one default, but completes as in (4a), obeying a mutually exclusive default, it should be obvious that completion must be occurring here on the basis of a type of evidence other than that provided by contour relatability. In particular, mergeability can account for the way completion occurs in all of these examples. Example (4a) satisfies complete mergeability, so completes, whereas example (3b) does not, so does not complete.

11. The meaning of mergeability

If 'mergeability' were merely defined as the ability of partially occluded volumes to merge behind an occluder, the definition would be circular. In order for the term 'mergeability' to have meaning we must specify how partially occluded volumes extend into occluded space and what relationships must hold between volume extensions behind an occluder in order for the partially occluded volumes to be mergeable. It might be thought that the mergeability of volume extensions can be estimated by the relatability of the principal axes of partially occluded volumes. On this approximation, principal axes would have to intersect behind an occluder. Moreover, the likelihood that volume extensions would merge would fall off as the angle between the extensions of their principal axes deviated from 180°, in the spirit of Kellman and Shipley's (1991) definition of contour relatability. However, volume completion behind an occluder seems to be much more flexible than this sort of relatability constraint would allow. For example, the principal axes of most of the examples in Fig. 9 between the two arrows are not relatable and do not intersect. The top portion of Fig. 2(a) might not even have a well-defined principal axis, yet completes with the lower portion. Fortunately, we do not need to limit our conception of mergeability to one of intersecting principal axes in order to formalize the role played by good continuation in the volume domain. We might argue, for example, that as the angle between any two intersecting lines running inside and parallel to the surfaces of their respective virtual volume extensions approaches 180°, the percept of amodal completion will be strengthened, and as it approaches 0° it will be weakened. Or, since the strength of mergeability and amodal completion appears to be inherently probabilistic in nature, we can formalize the notion of good volume continuation using the notion of probability distributions or ‘clouds’. Such clouds may attenuate with distance behind an occluder and have an oriented core of highest probability. We can then apply some generalized relatability criterion to this oriented core.

As Figs. 12 and 13 make apparent, the shape of the occluded volume extensions depends on the spatial relationship between the two visible partially occluded volumes. Since the shape of a partially occluded volume’s extension into occluded space is not determined solely by the visible portions of that partially occluded volume, we cannot determine local rules for how volume extensions should continue into occluded space. Rather, we can say that a visible partially occluded volume is consistent with a multitude of possible volume extensions. The strongest or most probable ones may be ones that continue visible surfaces and insides along their visible trajectory into occluded space. But trajectories that involve curving away from this most likely trajectory may be realized in the volume completion process in a way that depends on the relative location of completing parts. Since it would be impossible for the visual system to explicitly represent an infinity of possible volume extensions for each partially occluded volume, it appears that a purely bottom-up approach based on inferring volume extensions from local cues is untenable. Volume extensions are generated in interaction with global information about the scene, including information about distal partially occluded volumes. We can assign a cost to trajectories that deviate from the path most consistent with visible surface or principal axis curvatures, but we must accept that the completion process is inherently global and probabilistic. Indeed, even the shape of the occluded regions of the amodally completed figures in Fig. 13 are not precisely perceived (Tse, 1999).

There does not appear to be a simple way of formalizing the notion of mergeability according to a set rule that is either met or not met. Contour relatability was a rule that unambiguously separated images into those that permit amodal completion and those that do not. It could be formulated in more or less straightforward mathematical terms because it described conditions that were to be met by measurable image entities. In contrast, volumes are not measurable in the image, and volume extensions and completed forms seem to be inherently ambiguous and probabilistic in nature. It is natural to want hard and fast ‘laws’ that describe visual processing, since we wish to describe that processing with maximal certainty. Unfortunately, no such laws seem possible here because completion percepts come in degrees, and amodal completion seems to be influenced by multiple factors (Tse, 1999; see also van Lier’s discussion of fuzzy completion in this issue). By comparison, most linguists and philosophers of language gave up the search for necessary and sufficient conditions that would define a word that specifies class membership a few decades ago, and instead accepted that boundaries between classes are inherently probabilistic and ill-defined (See, e.g., Wittgenstein’s (1958) discussion of ‘family resemblances’). For example, in Fig. 13, when the blobs are close together, as in the central case, completion seems unambiguous.

However, as the two partially occluded blobs are displaced from one another, the strength of perceived amodal completion falls off. At some arbitrarily large displacement all observers would probably agree that completion fails. But there is no fixed boundary between the class of images that permit amodal completion and the class that does not. The best we can hope for is a list of the factors that affect the perceived strength of amodal completion, and a weighting of their respective contributions.

12. Other factors affecting completion

Complete mergeability is a more general explanatory principle for amodal completion over surfaces and volumes than contour relatability, even if it is probabilistic and non-measurable in nature rather than precise and measurable. However, there are other factors that contribute to the perceived strength of amodal completion than mergeability or good volume continuation. For example, in Fig. 15(a) amodal completion may seem weaker than in Fig. 3(4a) or the central case of Fig. 9 because the surface properties of the two bars are different. Similarly, amodal completion may seem weaker in Fig. 15(b) than in Fig. 1(a). Yin, Kellman and Shipley (1997) have provided evidence that similarity of surface features makes a contribution to the perception of amodal completion. This finding may be due to the fact that

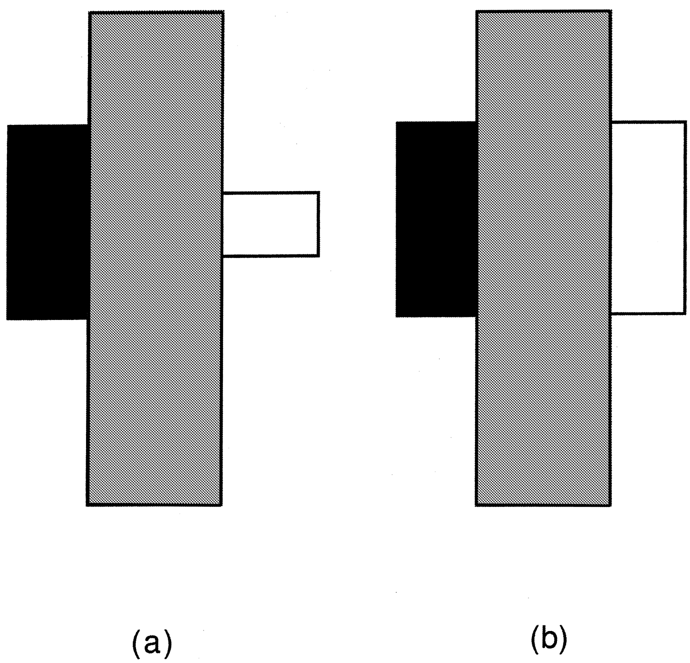

Fig. 15. Amodal completion is weakened if partially occluded surfaces are given different attributes.

partially occluded volumes are not just enclosures of space, but extended material. The visual system may be biased to link up partially occluded volumes of a common substance, and biased against linking up partially occluded volumes of a different perceived substance. Image regions with different features may entail surfaces with different properties, and surfaces with different properties are likely to be made of different material.

If this view is correct, then we would expect the visual system to determine amodal completion only once shape and material properties had been determined. If the visual system determines completion at the level of image properties, the rectangles behind the two cylinders in Fig. 16 should be judged to amodally complete with approximately equal strength since they have identical image properties. However, if the shadow is discounted before or during completion, then the rectangle in partial

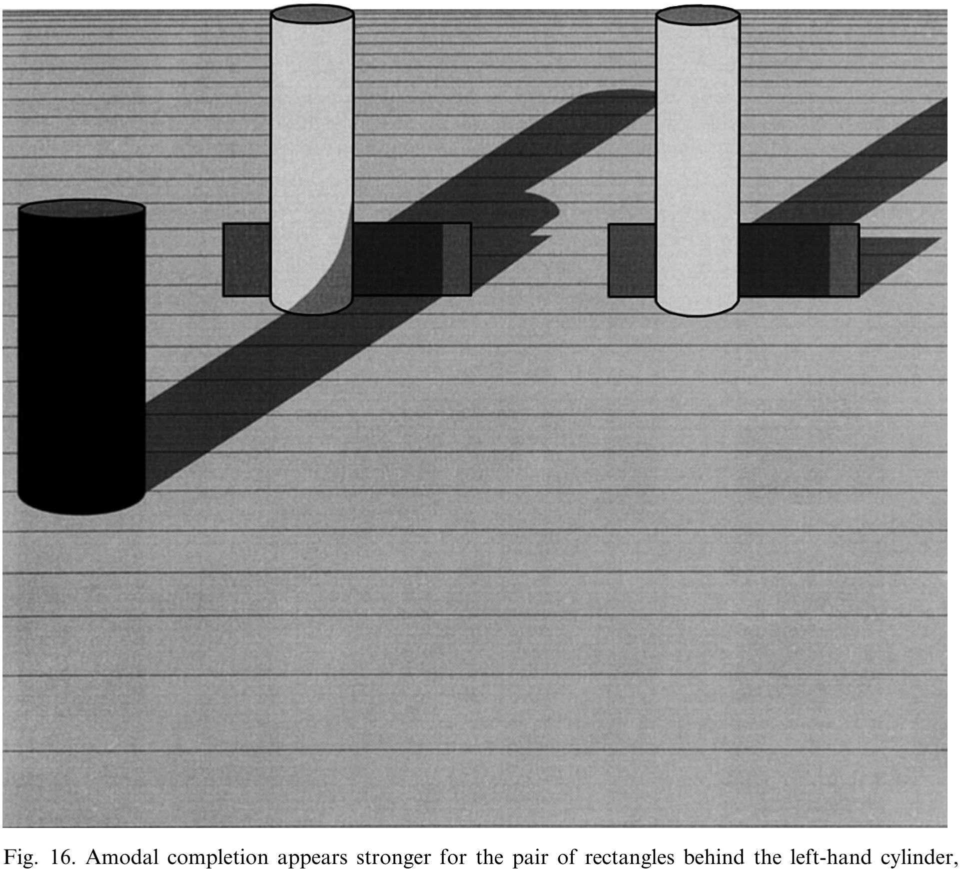

Fig. 16. Amodal completion appears stronger for the pair of rectangles behind the left-hand cylinder, suggesting that completion occurs only once surface properties have been determined based on global factors. In this case shadow information was discounted, allowing the left and right rectangles to have the same reflectance. This is not true for the rectangles in the right-hand case, since there is no shadow here, and color differences must be due to differences in reflectance.

shadow might have the same reflectance as the rectangle on the left side of the occluder. In that case, amodal completion should be stronger for the left-hand cylinder's case. In a similar drawing ten observers were asked which pair of rectangles appeared more like a unified object. Of these ten, all reported that the left-hand case appeared stronger. While more formal experiments need to be done, this again suggests that completion occurs at a level of information inferred from the image rather than over image information directly. Indeed, this result suggests that volume completion occurs after or in concert with a 'post-constancy' level of processing where shadows, illumination, and perhaps image size information has been discounted in order to recover the 'true' structure of the world that has generated an image.

In addition to surface features, the shape of contours clearly influences our percept of substance. Straight contours with sharp corners may tend to look like a harder material than smoothly curving contours because, in general, only a hard substance can support such angularity despite opposing natural forces, such as gravity. Thus, if the contours of the lower portion of Fig. 2(a) are altered so that it no longer appears to be a flowing gel, as in Fig. 17, the strength of the amodal completion percept may be weakened. Again, of the same ten observers, all reported that amodal completion was stronger in Fig. 2(a) than Fig. 17.

Several authors have challenged local-cue-driven accounts of completion like those of Kellman and Shipley (1991) or Wouterlood and Boselie (1992) and tried instead to explain completion in terms of global regularities in the patterns of

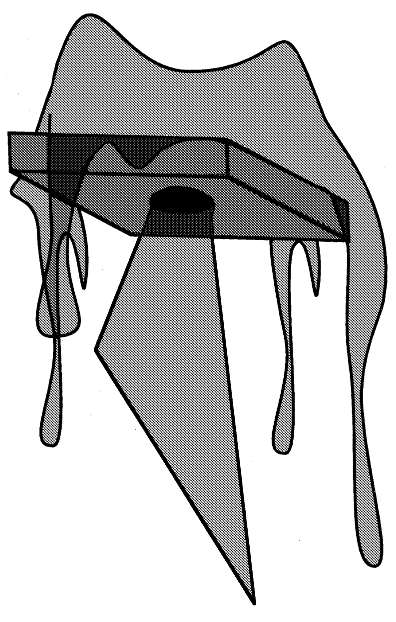

Fig. 17. The strength of amodal completion is weakened if the contours of the lower portion of Fig. 2(a) are modified so that it and the upper portion appear to be made of different substances.

completing objects (e.g., Buffart et al., 1983; Moravec & Beck, 1986; Sekuler, 1994; Sekuler et al, 1994; Van Lier et al., 1994, 1995). While global regularities may play a role in the decision whether partially occluded objects complete, amodal completion can take place in the absence of relatable contours, relatable surfaces, or global regularities, as in Fig. 2(a). This example completes because of complete mergeability. Nonetheless, the contribution of global geometric or pattern regularities (Van Lier et al., 1994; see Van Lier (1999)) to the percept of amodal completion can probably not be reduced to or derived from complete mergeability. In Fig. 18(a) the cross does not appear to amodally complete behind the gray square (into a small black square with two long 'appendages' emanating from its upper left corner) despite relatable contours and mergeable interiors. However, in Fig. 18(d), the black region appears to amodally complete more strongly behind the gray square. There are two possible reasons why the cross (Fig. 18(a)) may complete behind the occluder less strongly than the case in Fig. 17(d). First, the cross is symmetric. If the symmetry of the cross in Fig. 18(a) 'blocks' amodal completion, this would tell us that symmetry is one type of regularity in the image that the visual system seeks out and uses to generate a percept of objects in space. The second reason may be 'parallelism'. According to Lowe (1987) and Albert (1993) the closer that two parallel image contours are, the more likely that those contours will be regarded as intrinsic to (i.e. as projections from the edges of) the inferred surface that projects the region lying between the two parallel contours. It seems that parallelism may be more important than symmetry in blocking amodal completion here, because in Fig. 18(b), amodal completion is blocked in the absence of symmetry. Perhaps symmetry and parallelism both contribute to the blocking of amodal completion on the basis of mergeability here.

In general, 'vision abhors a coincidence'. According to the 'non-accidental viewpoint assumption' (Barrow & Tenenbaum, 1981; Binford, 1981; Lowe, 1985; Nakayama & Shimojo, 1992; Richards, Koenderink & Hoffman, 1987; Rock, 1984; see also Feldman, 1999) the visual system operates as if it were making the assumption that it is not perceiving an object from one of the few 'accidental'

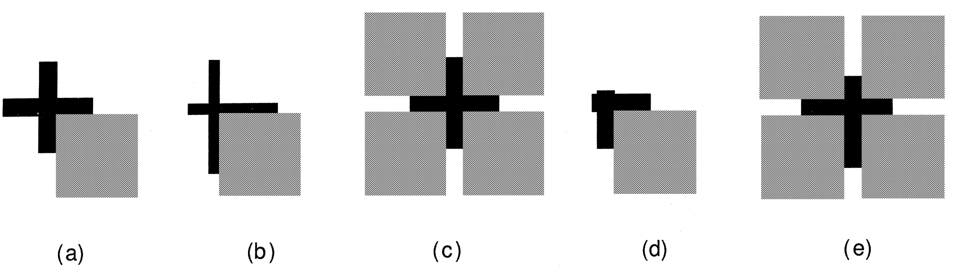

Fig. 18. In (a) observers tend to see a cross abutting a square, whereas in (d) amodal completion seems to take place. In (c) the cross interpretation may 'beat' the amodal completion interpretation, whereas in (e) the opposite may be true. To the extent that amodal completion fails in (b) it may be because of parallelism rather than symmetry.

viewpoints from which an object's surface layout is not derivable from its projected contours and other image cues. But an image is not solely a function of viewpoint. It is also a function of the arrangement of objects in the world. It would therefore be more accurate to speak of 'accidental images' rather than accidental viewpoints. An accidental image may be due to either an accidental viewpoint or a coincidental arrangement of objects in the world. But what counts as a coincidental arrangement for the visual system depends entirely on the types of regularities in the world that the visual system is looking for. In Fig. 18(a) the visual system must choose between two coincidences. Either the edges of the black cross exactly coincide with the edges of the gray square, or, if the black region amodally completes, then the appearance of a symmetrical cross or parallelism in the image is itself a coincidence. It appears that the coincidence of coinciding edges is less 'abhorred' than the coincidence of symmetry or parallelism, so the abutting cross interpretation 'wins' for Fig. 18(a). However, for Fig. 18(d), the amodal completion interpretation may win because the competing abutment interpretation is weakened by the lack of symmetry. Similarly, in Fig. 18(c), there may be a tendency to see a cross, but when symmetry is broken, as in Fig. 18(e), an amodally completing square interpretation may tend to dominate. By comparing the relative strength of image cues to object formation using a competition paradigm such as this, it may be possible to determine the relative importance of the various image regularities sought out by the visual system. Since there are many types of regularity in an image, it should not be surprising that many factors affect amodal completion, and that these factors can come into conflict. To the extent that they do conflict, the percept may tend to be ambiguous or 'flippable' between different interpretations, depending on mental set.

In Fig. 19(a), suggested by Rob van Lier, the global pattern appears to make a contribution to the percept of amodal completion. Note that here the partially

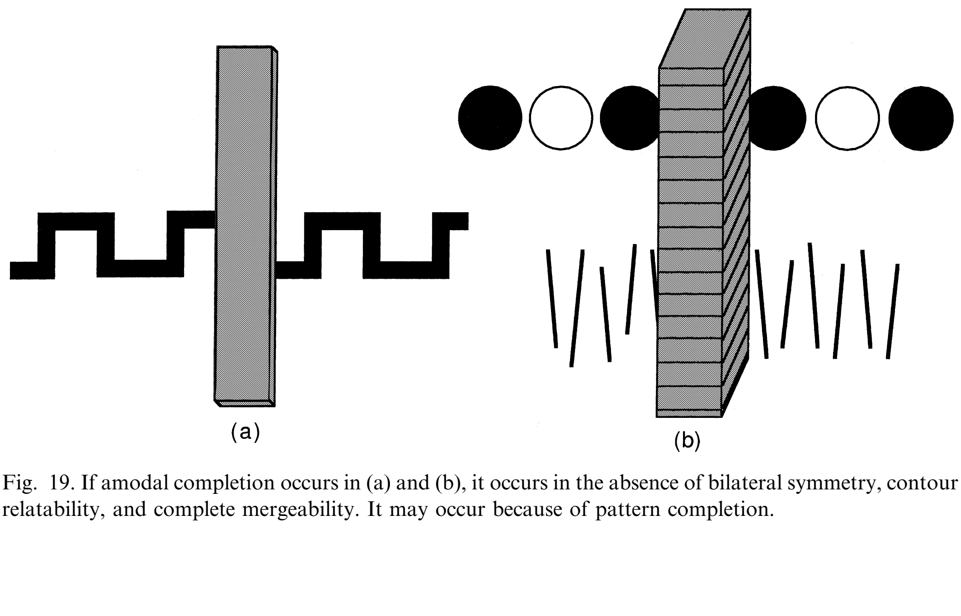

Fig. 19. If amodal completion occurs in (a) and (b), it occurs in the absence of bilateral symmetry, contour relatability, and complete mergeability. It may occur because of pattern completion.

occluded black zigzagged bars satisfy neither contour relatability nor complete mergeability, yet may appear to complete for some observers. Similarly, in Fig. 19(b) the (upper) series of circles or discs on either side of the occluder appears to amodally complete, as does the (lower) pattern of lines on either side of the occluder. Here what is merging may not be volumes or surfaces, but patterns. Thus, pattern regularities other than the regularity of bilateral symmetry may contribute to the percept of amodal completion, and may reflect the outcome of grouping rather than completion procedures.

Whereas Fig. 19 demonstrated that pattern completion processes can contradict edge, surface, and volume cues to completion, other examples demonstrate the opposite. Kanizsa (1979), (1955) gave examples where low-level cues to amodal completion appear to dominate the tendency to complete patterns. Kanizsa's insight can be extended to the volume domain, as in Fig. 20, where the pattern of short worms or separate blobs is overridden by stimulus-driven cues to amodal completion over partially occluded volumes.

This suggests that commonality or difference (of form, surfaces, edges, texture, color, luminance, depth, pattern and substance) may be the most fundamental

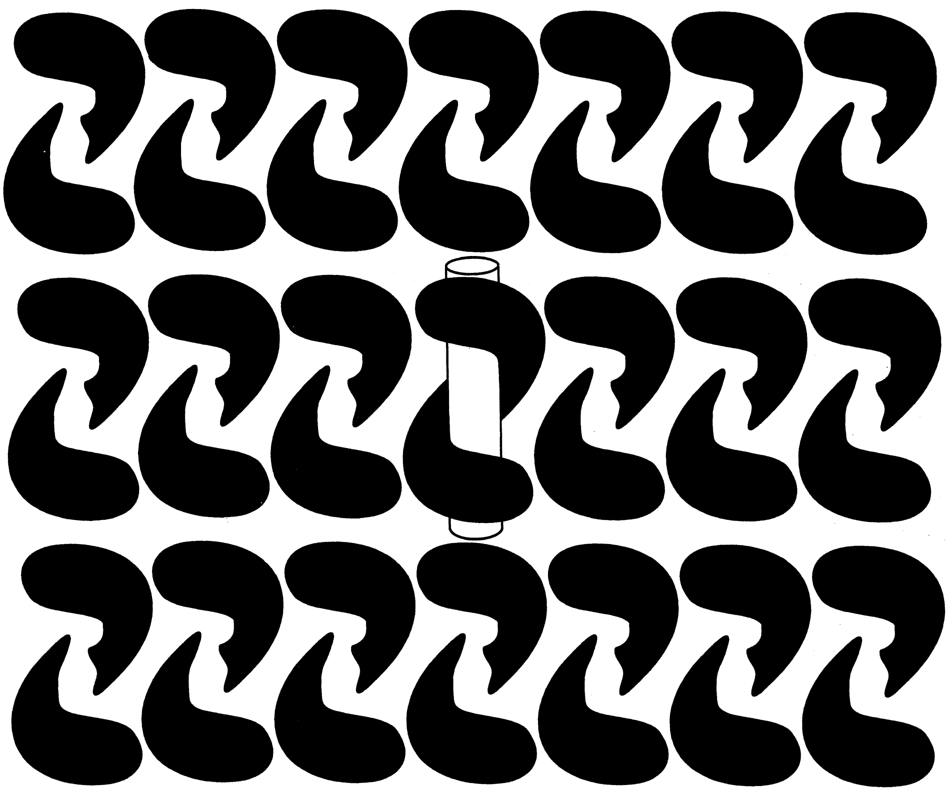

Fig. 20. Volume completion overrides the tendency to complete patterns.

factor determining whether amodal completion takes place or not. Since commonality can be determined at various levels (e.g. at the level of color, contour or volume analysis) it is possible that amodal completion is influenced by similarity decisions at multiple levels. Context may also play a role at various levels. For example, if we were to come across two nearby objects in a desert we might assume that they had something to do with each other simply because there was nothing else around them that could account for the improbability of their proximity. In a room full of objects, however, we might assume that they shared no special relationship, even when they shared the same spatial relationship as they did in the desert context. The ultimate strength of the amodal completion percept may depend on some interaction among multiple factors and levels of analysis that estimate the probability of a causal or physical relationship among the objects in the world that have generated inherently ambiguous image cues. However, it must be noted that the visual system cannot know a priori what should be completed as, say, just a pattern and what should be completed as a volume. Only by analyzing all cases of potential amodal completion as potential cases of pattern, surface, and volume completion, can some cases complete as volumes and others as patterns. Although the patterns in Fig. 19(b) may not complete as surfaces or volumes, this does not mean that they were never considered at a level of surface and volume analysis. As Fig. 20 makes evident, volume completion tends to dominate pattern completion, suggesting that pattern completion takes place only when no adequate 3D object interpretation is available. When amodal completion takes place over partially occluded surfaces or volumes, rather than patterns, complete mergeability must be satisfied.

13. General discussion

The visual system has evolved to capture biologically relevant information about the structure of the world. In order to accomplish this end, visual processing must be constructive in nature. The visual system overcomes the inherent ambiguity of image information by completing ambiguous or missing information so that completed information reflects the most likely structure of the world. One of the fundamental classes of image ambiguity that must be overcome in order to determine the shape and layout of objects in the world is ambiguity due to occlusion and self-occlusion. Since occlusion happens among 3D objects, it behooves the visual system to analyze occlusion relationships over internal representations of volumes rather than just contours, edges or visible surfaces.

Type (1) theories of amodal completion suggested that a percept of amodal completion is triggered by the existence of local image cues. For example, Kellman and Shipley (1991) and Wouterlood and Boselie (1992) argued that amodal completion follows from good contour continuation. Contour relatability, in particular, was an attractive basis for a theory of completion precisely because contour relatability is a measurable property of contours in the image. However, contour relatability is neither necessary nor sufficient for amodal completion to take place (see discussion of Fig. 2).

Indeed, there does not appear to be any image cue that is necessary or sufficient for the perception of amodal completion. Type (1) theories of completion are therefore inadequate. Type (2) theories emphasized the role of global pattern factors in completion, such as symmetry and regularity. Evidence described here clearly indicates that pattern-sensitive processes contribute to the perceived strength of amodal completion. Type (3) theories of amodal completion have insisted that image cues are first used to generate internal representations of surfaces (e.g. Nakayama et al., 1995) and volumes (Tse, 1998, 1999; Tse & Albert, 1998), and that amodal completion takes place only at the level where relationships among these internal representations are analyzed.

According to the type (3) account of completion developed here, the visual system uses both (1) local and (2) global image cues to generate representations of (3) volumes that may occlude one another under constraint of (2) learned or innate priors that minimize output complexity. According to this view, while the contour relationships emphasized by type (1) theories and global pattern regularities emphasized by type (2) theories play an important role in amodal completion, they are not direct cues to amodal completion. Rather, they and other image cues are used to infer edge, surface, and volume relationships in the world, and completion occurs at a level where these relationships are analyzed. In particular, it was postulated that the volume extensions of partially occluded volumes must be able to completely merge in order to complete into a larger volume. Even when a pattern cannot be completed as a volume, it is likely that it at least gets processed as being potentially a volume, since the visual system cannot know a priori what to complete as a pattern and what to complete as a volume. Moreover, volume completion tends to dominate pattern completion processes (Fig. 20). Thus, amodal completion over volumetric representations appears to be a central process in the construction of our perception of the 3D world.

In conclusion, vision involves construction of intermediate internal representations on the basis of global image information. A basic representational format of intermediate vision appears to be volumetric form. Higher-level processes, such as those involved in the determination of occlusion relationships, occur over these intermediate representations rather than directly over image cues. Completion processes are inherently probabilistic in nature and are influenced by multiple factors, including proximity, global context, pattern, similarity of material and form, relative orientation, and good volume continuation or mergeability.

Acknowledgements

I would like to thank Johan Wagemans, and the reviewers, Nicola Bruno and Rob van Lier for their very helpful comments on earlier drafts. In particular, Rob van Lier's review inspired Experiment 3. The author was supported by AASERT F49620-94-0376 to Patrick Cavanagh and McDonnell-Pew 98-49CNS.

References

Albert, M. K. (1993). Parallelism and the perception of illusory contours. *Perception*, *22*, 589–595.

Attneave, F., & Arnoult, M. D. (1956). The quantitative study of shape and pattern perception. *Psychological Bulletin*, *53*, 452–471.

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

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.

Feldman, J. (1999). The role of objects in perceptual grouping. *Acta Psychologica*, *102*, 137–163.

Helmholtz, H., von. (1962). *Treatise on physiologicalo optics* (Vol. 3). New York: Dover English translation by J P C Southall for the Optical Society of America, 1925, from the 3rd German edition of Handbuch der Physiologischen Optik, Hamburg, Voss, 1909; 1st ed. Leipzig, Voss, 1867.

Hochberg, J., & Brooks, V. (1960). The psychophysics of form: Reversible-perspective drawings of spatial objects. *American Journal of Psychology*, *73*, 337–354.

Hoffman, D. D., & Richards, W. A. (1984). Parts of recognition. *Cognition*, *18*, 65–96.

Hoffman, D., & Singh, M. (1997). Salience of visual parts. *Cognition*, *63*, 29–78.

Huffman, D. A. (1971). Impossible objects as nonsense sentences. *Machine Intelligence*, *5*, 295–323.

Kanizsa, G. (1979). *Organization in vision: Essays on gestalt perception*. New York: Praeger.

Kanizsa, G. (1955). *Riv. Psicolog.*, *49*, 7.

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

Lowe, D.G. (1985). *Perceptual organization and visual recognition.* Hingham, MA: Kluwer Academic Publishers.

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., Thinès, G., & Crabbé, G. (1964). Les compléments amodaux des structures perceptives. *Studia psychologica*. Louvain: Publications Universitaires.In G. Thinès, A. Costall, G. Butterworth, *Michotte's experimental phenomenology of perception*. Hillsdale, NJ: Lawrence Erlbaum, 1991.

Moravec, L., & Beck, J. (1986). Amodal completion: Simplicity is not the explanation. *Bulletin of the Psychonomic Society*, *24*, 269–272.

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 S.M. Kosslyn, D.N. Osherson, *Visual cognition: An invitation to cognitive science* (Vol. 2, 2nd ed., pp. 1–70). Cambridge: MIT Press.

Nakayama, K., Shimojo, S., & Silverman (1989). Stereoscopic depth: Its relation to image fragmentation, grouping, and the recognition of occluded objects. *Perception*, *18*, 55–68.

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*, 1168–1175.

Rock, I. (1984). *The logic of perception*. Cambridge MA: MIT Press.

Sekuler, A. (1994). Local and global minima in visual completion: Effects of symmetry and orientation. *Perception*, *23*(5), 529–545.

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

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), 977–993.

Tse, P. U. (1999). Volume completion. *Cognitive Psychology* (in press).

Tse, P. U., & Albert, M. K. (1998). Amodal completion in the absence of image tangent discontinuities. *Perception*, *27*, 455–464.

Van Lier, R. J. (1999). Investigating global effects in visual occlusion: From a partly occluded square to the back of a tree-trunk. *Acta Psychologica*, *102*, 203–220.

Van Lier, R. J., van der Helm, P. A, & 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., & 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.

Van Lier, R.J., & Wagemans, J. (1999). From images to objects: Global and local completions of self-occluded parts. *Journal of Experimental Psychology: Human Perception and Performance* (in press).

Waltz, D. (1975). Understanding line drawings of scenes with shadows. In P. H. Winston, *The Psychology of Computer Vision* (pp. 19–91). New York: McGraw-Hill.

Wertheimer, M. (1923). Untersuchungen zur Lehre von der Gestalt. *Psychologische Forschung*, *4*, 301–350.

Wittgenstein, L. (1958). *Philosophical investigations* (3rd ed., G. E. M. Anscombe, trans). New York: Macmillan.

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.