The occlusion illusion: Partial modal completion or apparent distance?

Stephen E Palmer, Joseph L Brooks, Kevin S Lai
Department of Psychology, Tolman Hall, University of California, Berkeley, Berkeley, CA 94720-1650, USA; e-mail: palmer@cogsci.berkeley.edu
Received 17 February 2006, in revised form 23 August 2006; published online 4 May 2007

Abstract. In the occlusion illusion, the visible portion of a partly occluded object (eg a semicircle partly hidden behind a rectangle) appears to be significantly larger than a physically identical region that is fully visible. This illusion may occur either because the visual system 'fills in' a thin strip along the occluded border (the partial-modal-completion hypothesis) or because the partly occluded object is perceived as farther away (the apparent-distance hypothesis). We measured the magnitude of the occlusion illusion psychophysically in several experiments to investigate its causes. The results of experiments 1 – 3 are consistent with the general proposal that the magnitude of the illusion varies with the strength of the evidence for occlusion, supporting the inference that it is due to occlusion. Experiment 4 provides a critical test between apparent-distance and partial-modal-completion explanations by determining whether the increase in apparent size of the occluded region results from a change in its perceived shape (due to the modal extension of the occluded shape along the occluding edge, as predicted by the partial-modal-completion hypothesis) or from a change in its perceived overall size (as predicted by the apparent-distance hypothesis). The results more strongly support the partial-modal-completion hypothesis.

1 Introduction

Although the visual system is remarkably accurate in representing the properties of most environmental objects under most circumstances, systematic illusions occur in the perception of many visual properties. Among the best known and understood of these are size illusions, most of which can be explained either by errors in the perceived distance to the target objects (eg the Ponzo illusion), by the influence of size contrast with nearby contextual objects (eg the Ebbinghaus illusion), or by size assimilation to nearby contextual figures (eg the Delboeuf illusion).

Perhaps the best known of all size illusions is the moon illusion (see figure 1a): the inaccurate perception of the moon as being larger when it is close to the horizon than when it is high in the sky. Although this phenomenon has been known for many centuries, its explanation is not fully settled (cf Hershenson 1989). The most widely accepted theory is based on differences in the apparent distance of the moon near the horizon versus high in the sky (eg Kaufman and Rock 1962; Rock and Kaufman 1962). The apparent-distance theory states that the horizon moon is (nonconsciously) perceived as farther away than the zenith moon—as though the sky were a 'flattened' dome—owing to the effects of distance cues, such as texture gradients, near the horizon. When the visual system takes this distance information into account, the seemingly more distant horizon moon looks larger than the seemingly closer zenith moon. A great deal of experimental evidence supports this explanation (Kaufman and Rock 1962; Rock and Kaufman 1962), which is a specific application of a more general perceptual regularity known as Emmert's law: If retinal size is held constant, then perceived size increases with increasing distance.

Another well-known type of explanation for size illusions is contextual contrast. A prototypical size-contrast effect is the Ebbinghaus illusion, in which two identical circles are perceived as different in size because of differential contrast with the size of

Figure 1. Examples of known size illusions including the occlusion illusion. (a) The moon illusion is an illusion of size putatively related to the misperception of distance. The moon appears larger on the horizon because it is perceived as farther away. (b) The Ebbinghaus illusion is thought to be a size-contrast illusion. The central circle surrounded by small dots appears to be larger than the central circle surrounded by large dots. (c) The Delboeuf illusion is an example of a size-assimilation illusion in which the size of a figure is distorted toward the size of nearby elements. The inner circle on the left and the outer circle on the right are identical. However, the one on the left looks larger because it is in the context of a larger circle. (d) The occlusion illusion may be a new type of illusion. The semicircle appears to be larger when adjacent to an occluder than when standing alone.

surrounding circles (see figure 1b). The central circle that is surrounded by a ring of many small circles appears larger than an identically sized circle surrounded by a ring of a few large circles. The explanation in terms of size contrast is that the perceived sizes of the central circles are influenced by the context provided by the surrounding circles, such that the small circles make the circle they surround appear larger, and the large circles make the circle they surround appear smaller.

A less well-known contextual explanation for size illusions is assimilation, exemplified by the Delboeuf (1892) illusion shown in figure 1c. Here, the inner circle on the left is objectively the same size as the outer circle on the right. Perceptually, however, the inner circle on the left appears larger than the outer circle on the right, ostensibly because these two circles in some way are assimilated with their context, seeming larger in the presence of the larger surrounding circle and smaller in the presence of the smaller surrounded circle. The conditions under which assimilation rather than contrast governs perceived size are not well understood.

In the present article we examine a size illusion, which Palmer (1999) called the 'occlusion illusion', that occurs when a single retinal region that is perceived as partly occluded (henceforth the 'target') appears to be larger than a physically identical region that is perceived as fully visible against a homogeneous background (the 'standard') (Kanizsa 1979; Kanizsa and Luccio 1978; Micali et al 1978; see Vezzani 1999 for a review). Figure 1d shows a canonical example of the occlusion illusion: two identical semicircles, with the target abutting a rectangle along its straight edge and the standard surrounded by a uniform background. The partly occluded target is generally seen as substantially larger than the fully visible standard. The question is: Why? The answer is of particular interest to us because we believe that it results from a fundamentally different

mechanism than other well-known size illusions: namely, partial modal completion that extends the visible portion of an object along a direction perpendicular to its occluded edge. This would be an explanation of a size illusion that relies on quite different perceptual principles than apparent distance, size contrast, or size assimilation. (We note, for the record, that the explanation in terms of partial modal completion implies that a subtle illusion of shape is also present, as we show in experiment 4, but the primary impression is of a difference in size.)

Crucially, the occlusion illusion may be just another example of a size illusion caused by errors in perceived distance arising from the fact that the target is perceived as *behind* the rectangle. If the standard is seen as lying in the same depth plane as the rectangle, whereas the target is behind the rectangle, then the target must be farther away than the standard and, by Emmert's law, correspondingly larger. This explanation is satisfying because it is plausible and appeals to a well known principle underlying other size illusions. Unfortunately, it does not square well with our phenomenological impression of the occlusion illusion. We believe the target appears larger by virtue of an extra strip of the circle being visible along its occluded edge.

It seems far less likely that the occlusion illusion can be explained by appealing to contextual contrast or assimilation. The fact that the occluding rectangle is larger than the target seems to contradict the size-contrast hypothesis, which predicts that the target should appear *smaller* than standard, the direction opposite to the actual illusion. To make a contrast explanation work, one must assume that the size of the standard is seen relative to that of its surrounding background, which is larger than the occluding rectangle, so that the target would seem larger than the fully visible one. Size assimilation has the reverse structure. It works if the rectangle operates as the context for the target and there is no appropriate context for the standard, but gives the wrong prediction if the surrounding background is taken to be the appropriate context for the standard. In either case, it is unclear why the illusion should be closely tied to partly occluded objects.

Our goal in this series of experiments is to better understand the nature of the occlusion illusion. The first three experiments explore the hypothesis that the illusion is indeed attributable to the perception of occlusion. In experiment 1 we measured the occlusion illusion psychophysically along with two variations in which we expected the illusion to disappear or be diminished, if it is indeed due to perceiving occlusion. The results also show that the illusion cannot be explained by size-contrast or size-assimilation effects. Experiment 2 rules out explanations in terms of local image structure (such as the presence of T-junctions) by showing that the effect can be produced by an illusory occluder in which there are no local differences immediately surrounding the target. Experiment 3 shows that the illusion is larger under conditions in which the evidence for occlusion is stronger, as determined by the shapes of the occluded and occluding objects. Finally, experiment 4 is a direct test whether apparent distance or partial modal completion provides the better explanation.

2 Experiment 1: Occluded, occluding, and surrounded conditions

In the first experiment, we used staircase psychophysical procedures to measure the standard version of the occlusion illusion and two control conditions. In the canonical 'occluded' condition, a larger rectangle abuts the target so that the target appears behind the rectangle (see figure 2a). In the 'occluding' condition, a smaller rectangle abuts the target so that the semicircle appears in front of the rectangle (see figure 2b). In the 'surrounded' condition, a larger rectangle surrounds the entire target (see figure 2c). If the illusion is due to size-contrast effects, it should be largest in the occluding condition, where the rectangle is smallest, and smallest in the surrounded condition, where the rectangle is largest. If it is due to size-assimilation effects, it should be

Figure 2. The displays for experiment 1. (a) The occluded condition. (b) The occluding condition. (c) The surrounded condition. (d) The occluding condition with the display-arrangement factor set such that the target is adjacent to the standard. (e) The occluding condition with the display-arrangement factor set such that the target is separated from the standard by the contextual rectangle. (f) The target at 3 degrees of occlusion by the contextual rectangle. The dotted line shows the extension of the bullet-shaped target behind the occluder. The left item shows the entire semi-circular portion of the target visible. The middle item shows a part of the semicircular portion occluded. This was the starting point for all of the staircases. The right item shows part of the bullet-like extension of the target.

largest for the surrounded condition, where the rectangle is largest, and smallest in the occluding condition, where the rectangle is smallest. If it is due to either partial modal completion or apparent distance, however, the illusion should be largest in the occluded condition and smallest in the occluding condition.

The portion of the circular region that was visible within the displays was varied from trial to trial. The unchanging standard to which the target was to be compared for size was an isolated portion of a circle that was entirely visible against the homogeneous gray background. A one-up/one-down staircase was used to find the point of subjective equality (PSE) between the visible portion of the target in the configural display and the simultaneously presented standard. The occlusion illusion is present if the target in the configural display that appears equal in size to the standard is actually smaller than the standard. The physical difference in size between the target and standard at the PSE can then be used as a measure of the magnitude of the illusion. If the illusion is due to partial modal completion of the occluded region at an occluding edge, the occluded condition should produce the largest illusion, because the straight edge of the semicircle is perceived as part of an occluding edge that belongs to the occluding rectangle, and the occluding condition should produce no illusion at all, because it belongs to the unoccluded semicircle. Predictions for the surrounded condition are less certain, because it can be seen either as a semicircle on top of the contextual rectangle or as a full circle that is partly inserted into a slit in the contextual rectangle, as many subjects indeed indicated when they were asked at the end of the experiment. It is reasonable to suppose that the surrounded condition might therefore produce an intermediate effect.

The apparent-distance hypothesis makes a similar prediction for the occluded condition, but is less clear in its implications for the occluding and surrounded conditions. One could argue that because the target in the occluding condition is unambiguously perceived as positioned in front of the rectangle, it should actually show a reversed illusion by being seen as smaller than the fully visible semicircle. This prediction is not firm, however, because one could equally well argue that the target is perceived as lying in the same depth plane as the standard (ie the picture plane) and the smaller rectangle as lying behind it. Predictions for the surrounded condition are even less clear because of ambiguity in relative depth.[1]

2.1 Method

2.1.1 *Subjects.* All thirteen participants were students at the University of California, Berkeley, who received partial course credit in their undergraduate psychology course. All gave informed consent, and the University of California, Berkeley, Committee for the Protection of Human Subjects approved the experimental protocol. The mean age of the participants was about 20 years.

2.1.2 *Design.* There were 24 display conditions resulting from the orthogonal combination of the following four factors: occlusion condition (whether the target was partially occluded, occluding, or surrounded), color of the target (black or white with the contextual rectangle being the opposite color), position of standard (left or right), and display arrangement (whether the standard was adjacent to the target of the configural display, figure 2d, or separated by the contextual rectangle, figure 2e). Notice that the display arrangement factor is confounded with distance, in that the standard was closer to the target in the adjacent condition than it was in the separated condition. Notice also that in the surrounded condition some of the contextual rectangle separated the target from the standard in both arrangements; the amount was simply

---

[1] We are currently pursuing experiments to more clearly indicate the depth relations between the contextual rectangle and the semicircle using stereoscopic cues to depth (Palmer and Schloss, in preparation). These manipulations should overcome such difficulties.

larger in the separated arrangement. Each of these 24 conditions defined a separate staircase procedure that ran until it converged, providing one PSE measurement for each condition for each subject.

2.1.3 *Displays and procedure.* Participants viewed the computer screen from approximately 50 cm. The size of the display was 14 inches diagonally and the resolution of the display was 1024×768 pixels at 60 Hz. Each display was presented on a neutral gray background. The three configural displays were as shown in figure 2. The standard was a half-circle of radius 2.24 deg (70 pixels) in all conditions and was always centered 7.97 deg (250 pixels) from fixation. The direction of its location (left or right from fixation) was determined by the position-of-standard factor. All examples in figures 2a – 2e have the position-of-standard factor set to left. The target was exactly the same as the standard in its semicircular part. Beyond the straight edge of the semicircular portion of the target, however, the figure was extended along horizontal tangent lines to create a bullet-shaped figure (as illustrated with the dotted line in figure 1f). In most cases, this extended portion would not have been visible because it was obscured by either the visible rectangular occluder (in the occluded condition as shown by the left item in figure 1f) or by an invisible occluder in the other two occlusion conditions (ie an occluder that is the same color as the background). This was done to ensure that, if more than the semicircular part of the target were made visible by the staircase procedure (as shown in the right item of figure 1f), no gap would appear between it and the contextual rectangle. The size of the visible portion of the target was varied by the staircase procedure according to the subject's responses. At the start of the staircase procedure in each condition, 1.60 deg (50 pixels) of the radius of the target was visible (as shown in the middle item of figure 1f).

The other parameters of the display differed for the three occlusion conditions. Each is described in turn. *Occluded condition (figure 2a)*: The occluding rectangle in the occluded condition was 5.45 deg (171 pixels) vertically by 4.43 deg (139 pixels) horizontally. Its distance from fixation depended on the display-arrangement factor. When the display-arrangement factor placed the rectangle between the standard and the target (separated), the rectangle was centered 5.74 deg (180 pixels) from fixation and directly adjacent to the straight edge of the target. For conditions in which the target and standard were adjacent (as in figure 2a) according to the display-arrangement factor, the rectangle was centered at 10.19 deg (320 pixels) from fixation directly adjacent to the straight edge of the target. The target was adjacent to the rectangle on either the left or right side depending on the combination of the display-arrangement and position-of-standard factors. For instance, if the standard appeared on the left side and the display arrangement specified the standard and target to be adjacent, then the target was on the left side of the rectangle as in figure 2a. The target was centered along the vertical extent of the rectangle. The horizontal position of the target varied depending on the state of the staircase. *Occluding condition (figure 2b)*: The rectangle in the occluding condition was 2.81 deg (88 pixels) vertically by 1.97 deg (62 pixels) horizontally. This rectangle abutted the straight edge of the target. The distance of the rectangle from fixation depended on the display-arrangement factor. When the rectangle separated the standard and target, the rectangle was 6.98 deg (219 pixels) from fixation. Otherwise, as in figure 2b, the rectangle was centered 8.95 deg (281 pixels) from fixation. The direction of this displacement from fixation depended on the position-of-standard factor. All other parameters were the same as in the occluded condition. *Surrounded condition (figure 2c)*: The rectangle in the surrounded condition was created by reproducing the displays from the occluded condition and simply adding a rectangle of the same size as the original rectangle in the occluded condition behind the target and abutting the original rectangle. The resulting larger rectangle was

4.44 deg (139 pixels) vertically by 5.42 deg (170 pixels) horizontally. All of the other parameters of the displays were the same as in the occluded condition. Notice that in this condition, the display-arrangement factor is better stated as closer versus farther rather than separated versus adjacent.

Subjects were told to look at each display and to indicate whether the target in the configural display was larger or smaller than the standard by pressing one of two buttons. In the initial display, the target was always much smaller than the standard. If the subject indicated that the target appeared smaller, then the next trial of that type showed a larger target by revealing an additional, single-pixel column of the partly occluded bullet. If the subject indicated that the target was larger, then the target on the next trial of that type was decreased in size by occluding an additional, single-pixel column of the partly occluded bullet. Each pixel column was approximately 0.03 deg wide. This procedure was continued until the subject had reversed his or her direction 8 times, and the PSE for that condition was computed as the average values from the last 6 reversals of the sequence for that display condition. This one-up/one-down staircase procedure (Levitt 1971) was followed separately for each of the 24 conditions described above, with the trials from the 24 staircases randomly interleaved. On each trial, the display remained on the screen until the subject responded. There was a delay of 500 ms between response and the onset of the next display. The program and displays of the experiment have been archived in the *Journal of Neurobehavioral Experiments* and can be downloaded from http://www.neuroexpt.com/ex_files/expt_view?id=157.

2.2 Results and discussion

The average PSEs across subjects are plotted in figure 3 for the three occlusion conditions when the targets were black versus white. An overall analysis of variance indicated significant main effects due to occlusion ($F_{2,24} = 35.11$, $p < 0.001$), color ($F_{1,12} = 5.18$, $p < 0.04$), and position ($F_{1,12} = 8.76$, $p < 0.01$), and significant interactions between occlusion and color ($F_{2,24} = 9.76$, $p < 0.001$), and occlusion and arrangement ($F_{2,24} = 9.23$, $p < 0.001$). No other main effects or interactions reached statistical significance.
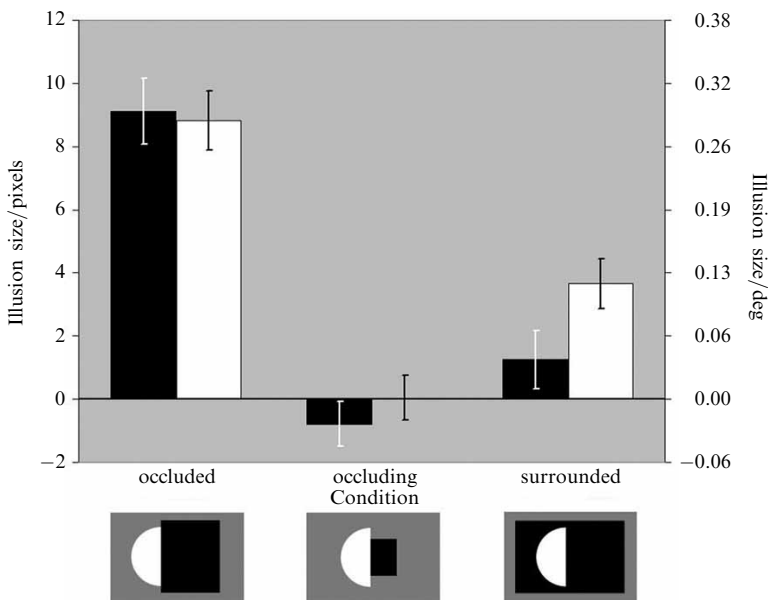
Figure 3. The results of experiment 1. Illusion size is shown in both pixels and degrees of visual angle. The color of the bars represents the color of the target for that condition.

The occluded condition produced a robust and highly reliable illusion in the expected direction: The target appeared to be the same size as the standard when about 9 fewer columns of pixels were visible in the target ($t_{12} = 9.31$, $p < 0.001$), which is about 20% smaller in area than the standard. The occluding condition produced no reliable difference ($t_{12} = -0.55$, $p > 0.50$). The surrounded condition was intermediate, producing a small, but reliable, illusion that averaged about 2 pixel columns (about 4% smaller) ($t_{12} = 3.17$, $p < 0.01$), which was significantly smaller than the illusion in the occluded condition ($F_{1, 12} = 29.06$, $p < 0.001$).

This pattern of results categorically rules out any explanation in terms of size-contrast effects. This hypothesis predicts that the largest effect should be found in the occluding condition, in which the contextual rectangle was smallest, but no illusion at all was observed in this condition. The pattern of results was also not consistent with an explanation in terms of size assimilation, which predicts that the size of the illusion should correlate with the size of the contextual rectangle (ie the black rectangles in figures 2a – 2c). Specifically, the target should be perceived as larger when the contextual rectangle is larger. The contextual rectangle is largest in the surrounded condition, smallest in the occluding condition, and of an intermediate size in the occluded condition. The size of the illusion does not fit this pattern, however. Consistent with size assimilation, the illusion is smallest in the occluding condition, in which the rectangle is smallest. However, the largest illusion is observed in the occluded condition (with an intermediate size rectangle) rather than in the condition with the largest rectangle (ie the surrounded condition). Size assimilation therefore cannot explain the effects we observed.

The results are fully consistent with the partial-modal-completion hypothesis because the order of the conditions is just what it predicts: strongest in the occluded condition, zero in the occluding condition, and intermediate in the ambiguous surrounded condition. They are somewhat less consistent with the apparent-distance hypothesis, which can be interpreted as implying that the occluding and surrounded conditions should produce a reversed illusion. No illusion was found in the occluding condition and a small illusion was present in the surrounded condition. Nevertheless, one could argue that when the target is seen as in front of the rectangle, it is perceived as lying in the same plane as the standard, in which case no illusion would be predicted. The results, therefore, are consistent with the apparent-distance hypothesis under the assumption that the closest figure in the configuration is perceived as lying at the same distance as the standard. In any case, it appears that the pattern of results is consistent with what would be expected if the size of the illusion were determined by the strength of the evidence for occlusion of the target in the configural displays.

Several other variables produced reliable effects in the data that appear to be due to factors other than occlusion. The illusion was slightly larger in the conditions in which the targets were white than when they were black ($F_{1, 12} = 5.18$, $p < 0.05$), but this was entirely due to the color effect in the surrounding condition. This effect is consistent with the widely-known irradiation illusion (von Helmholtz 1867/1962) in which a white square surrounded by a black frame looks larger than a black square of the same size surrounded by a white frame. It is unclear why the illusion was greater when the standard was on the left (5.05 pixels) than on the right (2.31 pixels) or why it was larger in the occluded condition when the target and standard were adjacent (8.48 pixels) than when they were separated (5.42 pixels), but larger in the surrounded condition when the target and standard were separated (3.19 pixels) than when they were adjacent (1.71 pixels). We note these effects for the sake of completeness, but do not have any coherent hypotheses to explain them.

Our preferred account of the primary results for the three occlusion conditions is that they arose from the degree to which subjects perceived the target as occluded by

the contextual rectangle. That is, the illusion was largest in the occluded condition, where participants clearly perceived the target to be occluded by the rectangle in that condition, and smallest (ie zero) in the occluding condition, where participants did not perceive the target to be occluded by the rectangle. Another possibility, however, is that the results are entirely driven by low-level visual features that correlate with occlusion in our displays, such as the presence of T-junctions. For example, there are two T-junctions with stems that 'point' toward the (farther) semicircular region in the occluded condition, whereas there are two T-junctions with stems that point away from the (nearer) semicircular region in the occluding condition, and no T-junctions in the surrounded condition. Perhaps the illusion is due simply to the presence of these image-based features rather than to the actual perception of occlusion per se. Experiment 2 was undertaken to address such issues.

3 Experiment 2: Illusory occluders

One way to produce perceived occlusion without introducing T-junctions around the target regions of these displays is to make the occluder an illusory figure (Kanizsa 1979). Indeed, an illusory figure can partly occlude an object without producing *any* image features immediately around the target in the configural displays, including differences in luminance or contrast (see figure 3). In the second experiment we therefore measured the magnitude of the occlusion illusion produced by the standard configuration (figure 4a) and two corresponding conditions in which the occluder is defined by illusory contours: the 'circle inducer' condition in which the inducing regions are four notched circles (or pacmen) just at the corners of the rectangle (figure 4b), and the 'complex-inducer' condition in which there are a larger number of more irregular inducing elements that define the same rectangular occluder (figure 4d). To provide a control for the illusory-occluder conditions, we included a 'reversed-inducer' condition in which the pacmen were rotated 180° so that their notches did not align to form illusory contours (figure 4c), even though their edges on the side of the target were still aligned with the edge of the target. If perceived occlusion is the sole determinant of the occlusion illusion, then the two illusory-occluder displays should produce an illusion equal to that in the standard occluded condition. If T-junctions—or indeed any image-based features immediately surrounding the target—are solely responsible, then neither of the illusory-occluder conditions should produce any illusion. Intermediate results for the illusory occluders are, of course, possible, particularly if the key factor is the strength of the evidence of occlusion. In this case, we would expect that the standard condition would produce the largest illusion, followed by the complex inducer and the circle inducer (in that order), with no illusion in the reversed-inducer condition.

3.1 Method

3.1.1 *Subjects.* All ten participants were students at the University of California, Berkeley, who received partial course credit in their undergraduate psychology course. All gave informed consent, and the University of California, Berkeley, Committee for the Protection of Human Subjects approved the experimental protocol. The mean age of the participants was about 20 years.

3.1.2 *Design.* The complete experimental design consisted of the four occlusion conditions described above: solid occluder, illusory occluder with circle inducers, illusory occluder with complex inducers, and the control condition with reversed inducers. The targets in all displays were black on a gray background, whereas the solid rectangular occluder and all of the inducing elements were white. The standard was always in the lower right quadrant of the display and the configural condition in the upper left quadrant in order to avoid subjects attempting to align the standard and target in any way.

Figure 4. Displays for experiment 2. (a) The standard configuration condition containing a solid rectangle occluder. (b) The circle-inducers condition with an illusory-rectangle occluder. (c) The reversed-inducers condition with no rectangle occluder. (d) The complex-inducer condition with an illusory-rectangle occluder.

3.1.3 *Displays.* Participants viewed the computer screen from approximately 50 cm. The size of the display was 14 inches diagonally and the resolution of the display was 1400×1050 pixels at 60 Hz. Each display was presented on a neutral gray background. The occluding rectangle for the solid-occluder condition and the illusory rectangles in the other conditions were 4.66 deg (200 pixels) wide and 9.30 deg (400 pixels) tall. It was centered 6.99 deg (300 pixels) to the left of the vertical midline of the screen and 2.33 deg (100 pixels) above the horizontal midline of the screen. The target abutted the rectangle centrally along the right side. The inducers for the illusory contour with circle inducers and the condition with reversed inducers had a radius of 1.95 deg (84 pixels). They were always white in color. The radius of the standard and target was 1.39 deg (60 pixels). The standard and targets were always black in color. The standard was centered 2.33 deg (100 pixels) below the horizontal midline and 8.84 deg (380 pixels) to the right of the vertical midline.

3.1.4 *Procedure.* The procedure was the same as in experiment 1 except for the following differences. Only four staircases were interleaved in a single block. Each subject participated in two blocks containing the four staircases. There was a short break between the two blocks.

3.2 Results and discussion

The PSEs for the two blocks were averaged for each subject. The PSEs averaged over subjects are plotted in figure 5 for the four occlusion conditions. The visible-occluder condition again produced a robust and reliable illusion ($F_{1,9} = 28.12$, $p < 0.001$) that was greater than that for any other conditions ($p < 0.01$ in every case). The complex-inducer condition also produced an illusion that was reliably greater than zero ($F_{1,9} = 7.11$, $p < 0.05$) and reliably greater than the reversed-inducer condition ($p < 0.001$), but not reliably greater than the circle-inducer condition ($F_{1,9} = 4.26$, $p < 0.10$). The circle-inducer condition was not quite reliably greater than zero ($F_{1,9} = 3.38$, $p = 0.10$), but was reliably greater than the reversed-inducer condition ($F_{1,9} = 26.74$, $p < 0.01$). The reversed-inducer condition produced no illusion at all, giving an average PSE that was slightly in the wrong direction ($F_{1,9} = 1.12$, $p > 0.30$).
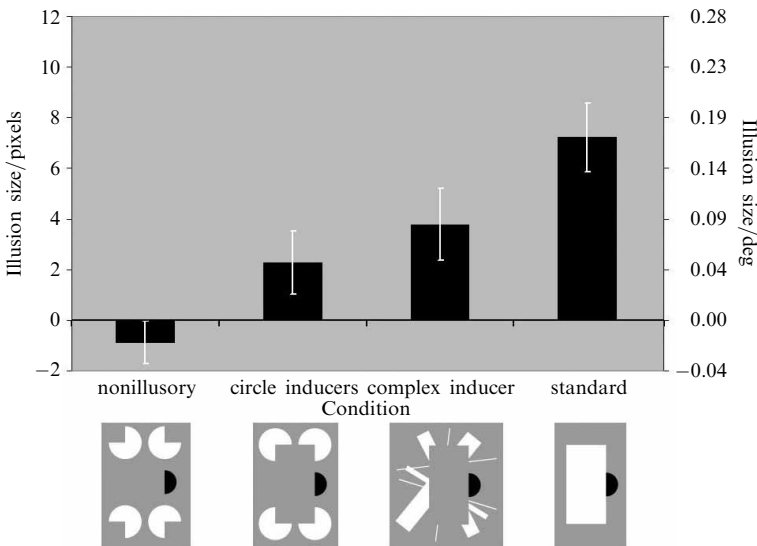
Figure 5. Results of experiment 2. Illusion size is shown in both pixels and degrees of visual angle. Notice that the scale for degrees differs from that of experiment 1 because the screen resolution was different in the two experiments.

The intermediate illusion effects in the illusory-contour conditions are ambiguous concerning the determinants of the illusion. These illusory occluders, with no direct cues to occlusion immediately surrounding the target, did produce size illusions in the predicted direction, thus ruling out any explanation solely in terms of the presence of T-junctions or other sorts of local structure around the target. However, the magnitude of the illusion was significantly less than that in the standard condition, which did contain T-junctions and related luminance structure consistent with the perception of occlusion. This result is consistent with findings in a previous study with illusory occluders (Perussia 1983; as reported in Vezzani 1999). This indicates that factors like T-junctions and related luminance structure are also relevant to producing the illusion.

Perhaps the most parsimonious description of the results is that stronger perceptual evidence for occlusion produces larger magnitudes of the illusion. This possibility might arise from probabilistic effects, quantitative effects, or both. The strictly probabilistic view is that whenever occlusion is perceived, an occlusion illusion of a fixed magnitude occurs, and when it is not perceived, no illusion occurs. When the evidence for occlusion is strong, as in the solid-occluder condition, subjects perceive occlusion (and thus experience the illusion) on a high proportion of trials, but when the evidence is weaker, as in the illusory-occluder conditions, they perceive occlusion (and experience

the illusion) on a smaller proportion of trials. Over the entire experiment, then, these probabilities would produce smaller illusory effects when the evidence for occlusion is weaker. The strictly quantitative view is that the strength of the illusion on any trial varies directly with the amount of sensory evidence favoring occlusion. These two views predict different distributions of illusion magnitudes over trials for displays with weaker evidence of occlusion (such as the illusory occluders), with the probabilistic view predicting bimodal distributions with larger variance, and the quantitative view predicting unimodal distributions with smaller variance. Unfortunately, the PSE data available from psychophysical staircase methods do not preserve such information, and so we cannot test these predictions.

4 Experiment 3: Effects of region shape

The third experiment was undertaken primarily to explore the possibility that the strength of the illusion could also be influenced by the strength of the evidence for occlusion via global shape considerations (see figure 6). We reasoned that a square partly occluded by a rectangle (figure 6a) would be more consistent with a non-occluding 'mosaic' interpretation (ie two adjacent rectangles that share a border in the same depth plane) than the corresponding display in which a circle is partly occluded by the same rectangle (figure 6b) even though they both contain T-junctions. Exactly why this might be the case is a deep and important problem that is beyond the scope of the present article. We believe that the reasons are related to the greater ecological likelihood that the pair of adjacent rectangles would arise from a scene without occlusion than would a rectangle adjacent to a semicircle, but we know of no statistical evidence supporting this conjecture.

(a)

(b)

(c)

(d)

Figure 6. The displays for experiment 3. (a) A rectangle occluder with a rectangle target. (b) A rectangle occluder with a circular target. (c) An oval occluder with a rectangle as the target. (d) An oval occluder with an oval as the target.

In any case, if the adjacent rectangles produce weaker evidence of occlusion than the rectangle adjacent to the target (for whatever reason), then we expect the pair of rectangles to produce a weaker illusion than the standard display. To be sure that any such effects were not simply due to the rectangular shape of the target region itself, we also included two other conditions in which the occluder was an ellipse (figures 6c and 6d), reasoning that in these cases both the partly occluded circle and the partly occluded rectangle were relatively unlikely to arise without involving occlusion. We therefore expected that they both would produce larger illusions than in the condition with a rectangle occluding another rectangle (figure 6a). We also collected data more directly relevant to the strength of the perceived evidence for occlusion by asking different participants to make explicit ratings about the degree to which the upper target region appeared to be behind and occluded by the larger, lower region for each of our conditions. If the magnitude of the illusion varies according to the strength of the perceptual evidence for occlusion, these ratings of depth and occlusion should correlate highly with the measured magnitude of the illusion.

Another issue addressed in this experiment is the possible effect of the orientation of the occluding edge. The reason that orientation is of interest is that one possible filling-in explanation of the standard phenomenon would be in terms of da Vinci stereopsis (Nakayama and Shimojo 1990). That is, because the target is perceived as occluded behind the rectangle (owing to monocular depth cues, in this case), the visual system may assume that a thin strip of the target along the border with the rectangle should be seen in one eye but not in the other. This thin strip of visual information must be integrated into the overall binocular percept at the edge where the depth difference has been registered. Notice, however, that the display for the occlusion illusion contains no actual depth difference at the target/rectangle edge and thus there is no disparity in the retinal images. Regardless, because of the perceived difference in depth, the visual system may assume that a thin strip should be filled in along the edge where one eye would have seen it. In this case it might be filled in from a 'da Vinci buffer' that is normally used for the portion of any partly occluded object that would be visible in only one eye, if the object were indeed partly occluded. If this were the case, however, the illusion should disappear, or be greatly diminished, if the occluding edge is horizontal, because then there would be no significant da Vinci stereopsis. This prediction was tested in experiment 3 by including conditions in which the occluding border was horizontal as well as ones in which it was vertical.

4.1 Method
4.1.1. *Subjects.* All twelve participants were students at the University of California, Berkeley, who received partial course credit in their undergraduate psychology course. All gave informed consent, and the University of California, Berkeley, Committee for the Protection of Human Subjects approved the experimental protocol. The mean age of the participants was about 20 years.

4.1.2 *Design.* The 16 display conditions were defined by the orthogonal combination of the following factors: occluded shape (circle or rectangle), occluding shape (oval or rectangle), occluding edge orientation (globally horizontal or vertical), and position (standard on the left side or right side). The standard was always at the top of the screen and the configural display on the bottom in a diagonal arrangement such that neither the horizontal nor the vertical dimensions were aligned.

4.1.3 *Displays.* Participants viewed the computer screen from approximately 50 cm. The size of the display was 14 inches diagonally and the resolution of the display was 1024×768 pixels at 60 Hz. Each display was presented on a neutral gray background. For all conditions, the occluder was black and the target and standard shapes were

both white. In all conditions with the rectangle as the occluding shape, the rectangle was 3.19 deg (100 pixels) by 5.46 deg (171 pixels). The oval occluder was 3.19 deg (100 pixels) on its shortest dimension and 6.38 deg (200 pixels) on its longest dimension. These occluders were centered 7.97 deg (250 pixels) from the vertical midline and 6.41 deg (201 pixels) from the horizontal midline when horizontally oriented. In the vertically oriented conditions, the occluder was centered 6.38 deg (200 pixels) from the vertical midline and 4.79 deg (150 pixels) from the horizontal midline. The orientation of the rectangle and oval occluders depended on the orientation factor. When the orientation factor was vertical, then the target abutted the right edge of the occluder. In horizontal-orientation conditions, the target abutted the top edge of the occluder. The partial rectangle was 3.99 deg (125 pixels) along its occluded edge. It was the same size in the other dimension. However, the amount of the other dimension that was visible varied with the staircase procedure. The radius of the target was 2.01 deg (63 pixels). It was extended into a bullet shape as in experiment 1. The amount of this shape that was visible also varied with the staircase procedure. The standard was a 2.01 deg (63 pixels) radius semicircle. The standard rectangle was 3.99 deg (125 pixels) by 2.01 deg (63 pixels). The standard was oriented in the same direction as the target. It was located the same distance from fixation as the target but in the diagonally opposite location on the screen.

4.1.4 *Procedure*. The procedure was the same as in experiment 2 except that there were 16 independent staircases being run in an interleaved fashion. Each subject completed only a single run of the procedure, so there was just one estimate of the PSE for each condition for each subject.

4.2 *Results and discussion*
The average PSEs over subjects are plotted in figure 7 for the four occlusion conditions. An overall analysis of variance showed a significant main effect of occluded shape ($F_{1,11} = 45.78$, $p < 0.001$) and a significant interaction between occluded shape and occluding shape ($F_{1,11} = 13.95$, $p < 0.01$). This is the interaction plotted in figure 7. No other factors or interactions were statistically reliable, including any that included the orientation of the occluding edge. da Vinci stereopsis, therefore, cannot be a significant factor in the explanation of the occlusion illusion.
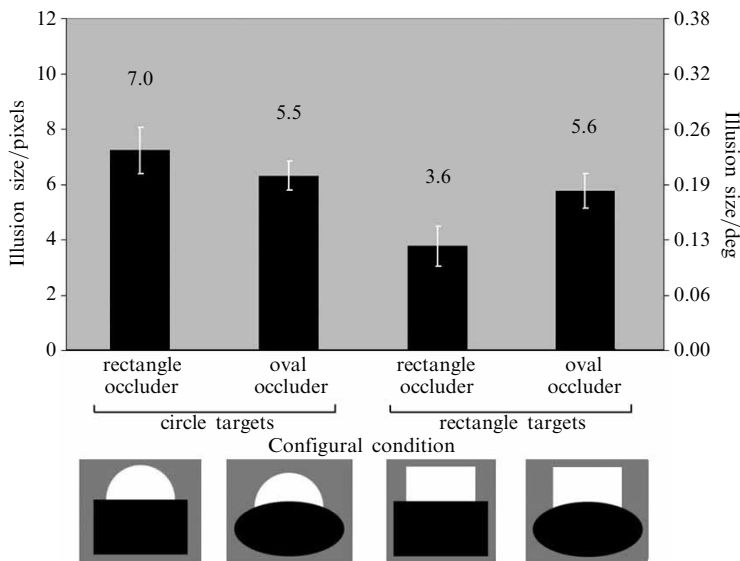
Figure 7. Results of experiment 3. Illusion size is shown in both pixels and degrees of visual angle.

The standard circle/rectangle condition again produced a robust and reliable illusion in which the PSE for the target circle was about 7 pixel columns (or 15%) smaller than the standard ($t_{11} = 8.63$, $p < 0.001$). The corresponding square/rectangle condition produced a significant illusion of about 4 pixel columns (or 8% smaller in area) ($t_{10} = 5.21$, $p < 0.001$), which was reliably smaller than the circle/rectangle condition ($F_{1, 11} = 46.11$, $p < 0.001$). The circle/ellipse condition and square/ellipse conditions also produced significant illusions of about 6 pixel columns (about 13% smaller in area) ($t_{11} = 12.21$ and 9.19, $p < 0.001$), which were reliably larger than the square/rectangle condition ($F_{1, 11} = 57.41$ and 11.95, $p < 0.01$ in both cases) but only slightly smaller than the size of the illusion in the circle/rectangle condition ($F_{1, 11} = 4.32$ and 1.58, $p > 0.05$ in both cases). The results were thus generally consistent with expectations based on our introspective intuitions about the extent to which the different conditions supported perception of occlusion.

To find out whether our intuitions were representative of those of other observers, we showed the four configural displays to a group of eight naive subjects, who had never seen the occlusion illusion and did not know of its existence. They were asked to rate each condition in terms of the degree to which the upper region appeared to be behind and occluded by the lower region on a 9-point scale (with 9 representing a strong perception that the lower region is closer and 1 representing a very weak perception that the lower region is closer). There were five repeated measures for each display, and median ratings were averaged across subjects.

The average ratings are given above the corresponding histograms in figure 7. Consistent with our intuitions, the circle/rectangle received the highest ratings and the square/rectangle received the lowest ratings, with the other two cases intermediate between them. Statistically, the circle/rectangle condition was rated significantly higher than both the circle/oval and square/rectangle conditions ($p < 0.05$ in both cases), the square/rectangle condition was rated significantly lower than any of the other conditions ($p < 0.01$ in each case), and the circle/oval and square/oval conditions did not differ significantly from each other ($p > 0.50$). Quantitatively, the magnitudes of the average occlusion ratings show a remarkably strong linear relationship with the psychophysical measurements in the magnitudes of the illusion effects in the same four conditions ($r = 0.97$, $p < 0.05$). This result is thus consistent with the hypothesis that the magnitude of the illusion follows the strength of the perceived evidence for occlusion.

5 Experiment 4: Perceived distance versus partial modal completion

The previous three experiments support the general conclusion that the magnitude of the occlusion illusion is determined by the strength of the perceptual evidence favoring occlusion, but they do not answer the question posed at the outset: namely, whether the occlusion illusion is better explained by perceived distance or by partial modal completion. Note that our displays thus far have implicitly been based on the assumption that the illusion is due to partial modal completion, because we varied the size of the target in the configural display by occluding a larger or smaller portion of a bullet-shaped region. We thus maintained a constant radius of the partial-circle portion of the bullet rather than increasing the overall size of the target by increasing its radius. In other words, we actually changed the *shape* of the target as well as its size rather than changing its size alone.

In the fourth experiment, we tackle the question whether the occlusion illusion is more consistent with the perceived-distance hypothesis or the partial-modal-completion hypothesis using the following two-part method. In the first phase of the experiment, we found the PSEs for each participant of a fixed, configural display using two different sets of variable-sized test figures. The 'shape-based' test figures were defined, as in the three previous experiments, by occluding more or less of the same-sized

bullet-shaped region. Notice that the radius of the test figures in this series does not change, but both their shape and overall size does. The 'size-based' set was defined by making the test figure proportionally larger or smaller overall (figure 8). Here, the radius changes, but the shape does not. After the PSEs were found for both the shape-based and size-based test figures for a given participant, he or she was shown the single configural display together with the two targets that the same participant had just judged to be the same size as the circular part of the unchanging configural display, one being the size match from the shape-based series and the other being the size match from the size-based series. The subject was then given the two-alternative forced-choice (2AFC) task of indicating which of the two test figures looked more like the circular part of the configural display. If the partial-modal-completion hypothesis is correct, subjects should systematically prefer the 'shape-based' figure that has the same radius as the target of the configural display but a different shape. If the apparent-distance hypothesis is correct, they should systematically prefer the 'size-based' figure that has the same shape as the target in the configural display but a different radius.

Figure 8. The standard shape changed in two ways in experiment 4. In the shape-change series, both the shape and size of the target was changed by moving it in and out from behind an occluder. In the size-change series, only the size of the target region was changed by dilating it.

5.1 Method

5.1.1 *Subjects.* All twelve participants were students at the University of California, Berkeley, who received partial course credit in their undergraduate psychology course. All gave informed consent, and the University of California, Berkeley, Committee for the Protection of Human Subjects approved the experimental protocol. The mean age of the participants was about 20 years.

5.1.2 *Design.* The experiment consisted of two parts. In the first part of the experiment four display conditions were defined by the orthogonal combination of the following factors: occluder and occluded-shape combination (circle with rectangle occluder or rectangle with oval occluder) and size of the occluded region (large or small). There were two staircases for each of these four conditions. One staircase varied the size of the standard by dilation. The other staircase varied the size and shape of the standard by varying the portion of the bullet-shape that was visible. (Note that, unlike experiments 1–3, the fully visible 'standard' region is the one that varies in the present experiment, and the partly occluded 'target' region is constant.)

5.1.3 *Displays.* Participants viewed the computer screen from approximately 50 cm. The size of the display was 14 inches diagonally and the resolution of the display was 1024×768 pixels at 60 Hz. Each display was presented on a neutral gray background. For all conditions, the occluder was black and the target and standard regions were white. The variable unoccluded standard shape was always located in the upper-left quadrant of the display at the same coordinates as in experiment 3. The size of the standard region was adjusted according to the staircase procedure. The unchanging configural display with the partial target shape and occluder were always located in the lower-right corner of the display, as in experiment 3. The rectangle and oval occluders in this experiment had the same dimension as those in experiment 3. The fully visible standards varied in size depending on the staircase procedure. The small semicircle had a radius of 1.43 deg (45 pixels). The large semicircle had a radius of 2.01 deg (63 pixels), which was the same as in previous experiments. The standard varied in two ways from these parameters. In the dilation-change staircase, the radius could vary up to 0.64 deg (20 pixels) in each direction. In the occlusion/disocclusion staircase, the standard could be occluded or disoccluded 0.64 deg (20 pixels) in each direction. The partially occluded rectangle was 2.87 deg (90 pixels) vertically and 1.40 deg (44 pixels) at its widest point in the small condition. It was 3.99 deg (125 pixels) vertically and 1.97 deg (62 pixels) at its widest point in the large condition. The standard rectangle on the other side of the screen varied in the same manner as the circle according to the staircase procedure. In the dilation staircases, the aspect ratio of the rectangle was maintained but the size of the horizontal and vertical dimensions of the rectangle was adjusted.

5.1.4 *Procedure.* The procedure was similar to that of the previous experiments except for the changes noted above regarding which part of the display changed with the staircase and the two different types of changes. At the end of the staircase procedures, there were 32 additional trials, in which the 2AFC task was run. In these trials, subjects were presented with one of the configural displays that they had seen repeatedly in the staircase trials together with both of the test figures (ie fully visible standards) that they themselves had judged to be the same in overall size as the corresponding region in the configural display. They were then required to make a forced-choice response indicating which of the two test figures looked more similar to the target in the configural display that was presented at the same time. They performed this forced-choice task for each of the four conditions defined above a total of 8 times per condition.

5.2 *Results and discussion*
Figure 9 shows the results of the 2AFC task averaged over subjects for the two shape-configurations conditions (circle occluded by rectangle and rectangle occluded by oval) and two size conditions (large and small) plotted in terms of the probability of participants choosing the shape-based alternative. Probabilities greater than 50% thus favor the partial-modal-occlusion hypothesis, whereas probabilities below 50% favor the apparent-distance hypothesis. The critical result is that, over all conditions, participants chose the shape-based alternative as looking more like the partly occluded shape than the size-based alternative on about 66% of the trials ($t_{11} = 4.14$, $p < 0.001$), and all four conditions exhibited effects in this direction. As figure 9 shows, the magnitude of this preference varied across both the size factor (large versus small versions) and the shape factor (circle versus rectangle occluded figures) in a roughly additive way. For the partly occluded circle, the shape-based figure was chosen 72% of the time, on average, which is significantly above chance (50%) ($t_{11} = 4.36$, $p < 0.001$). For the partly occluded rectangle, the shape-based figure was chosen 59% of the time, which again is above chance, but not significantly so ($t_{11} = 1.78$, ns). The same pattern was evident for the large versus

Figure 9. Results of experiment 4. The percentage of participants who chose the shape-change test figure for the large and small versions of the two shape configurations.

small versions, although, surprisingly, the preference was greater for the smaller than the larger versions ($t_{11} = 2.25$, $p < 0.04$). In general, the magnitude of the bias toward choosing the shape-based alternative across conditions covaried with the magnitude of the illusion—ie larger illusion sizes produced a greater preference for choosing the shape-based alternative—although with only four data points this trend was not statistically reliable ($r = 0.58$, ns). Such a finding is reasonable, given that as the size of the illusion diminishes to zero the 2AFC data should asymptote to chance.

6 General discussion

The results of the experiments reported above provide several important insights into the nature of the occlusion illusion. First, size contrast and size assimilation cannot be significant factors because the pattern of results in experiment 1 for rectangles of different sizes directly contradicts the predictions based on both of these factors. Second, the illusion depends on the perception of occlusion rather than just on the presence of explicit T-junctions or other local luminance structure, because experiment 2 showed that it can be obtained, albeit in attenuated form, when the occluder is an illusory figure. Third, the strength of the illusion appears to depend on the strength of the perceptual evidence for occlusion, consistent with the constellation of findings in experiments 1, 2, and 3. Fourth, the partial-modal-completion hypothesis provides a better explanation of the results than the apparent-distance hypothesis because participants in experiment 4 reliably chose the shape-based comparison figure over the size-based comparison figure as looking more similar to the partly occluded figure in the configural display, even when they had been matched by the same subject for perceived size.

We take the occlusion illusion to be a very general phenomenon that occurs whenever an object or surface is perceived as partly occluded by a shared edge. Partial occlusion is ubiquitous in normal, everyday visual perception, and there do not appear to be any unusual or special circumstances required to obtain the illusion that would restrict its generality. This implies that whenever the visual system amodally completes

an object behind an occluder—a very frequent occurrence—there is a small *modal* component to this completion. It is not generally noticed because it is relatively small (although we measured up to a 20% effect in some cases), and, more importantly, because it is *noticeable* only under the relatively unusual conditions when a physically identical unoccluded region is available for comparison. Nevertheless, we believe that it is present whenever there is visual perception of occlusion and amodal completion.

Before closing, we wish to discuss briefly three important issues that have not been answered by the results presented above. One is how quantitative variations in the strength of the illusion are properly understood. Do they result from probabilistic variations in consciously seeing the target figure as partly occluded versus not, or do they arise even when the observer always sees the target as partly occluded as a result of variations in the strength of the underlying sensory evidence for occlusion? Our own introspective experiences suggest the latter, because we have measured such variations in our own data when we always perceived the target figures as partly occluded, but we have not yet systematically collected trial-by-trial data from naive participants on this issue. We also plan to conduct experiments using stereoscopic displays that provide unambiguous information from binocular disparity which figures lie in front of versus behind each other to see if this eliminates quantitative differences.

A second issue is how to understand the relation between the 'one-sided' occlusion illusion we have studied above versus a 'two-sided' version of what seems to be the same situation (often called the 'shrinkage' illusion) that is illustrated in figure 10. Here the partly occluded figure is visible on both sides of the occluder. On the basis of the present findings with the one-sided illusion, one would expect that the partly occluded object in the two-sided version would appear to be *much* larger than the objectively same-sized disoccluded figure that is fully visible. In fact, the partly occluded figure is now seen as significantly *smaller* than its unoccluded version (Kanizsa 1979; Kanizsa and Luccio 1978; for a review of work on this version of the illusion see Vezzani 1999).

Figure 10. Two different illusions. One side of the partially occluded circle in (b) looks larger than an isolated target of the same size (a), consistent with data from experiments 1 – 3. However, the entire partially occluded circle, consisting of both separate regions (b) looks smaller overall than a full circle of the same size (c). This 'shrinkage' illusion is thus opposite of the occlusion illusion as studied in the present article.

One possibility is that there is a concomitant *decrease* in the apparent size of the invisible (but implicitly perceived) portion of the occluded object, which is not noticeable in the one-sided version, but which dominates in the two-sided version (Palmer 1999, pages 326–327). Another is that the shrinkage illusion is not due to occlusion at all, but to some more general extrapolation process. Clearly, understanding the relation between these two different forms of what seems superficially to be the same illusion is an interesting and important goal for future research.

The final, and in many ways most perplexing, issue is how to understand the nature of the partial-modal-completion hypothesis. Specifically, what does it mean for the visual system to 'fill in' a thin strip along the occluded edge? Where is this thin strip located, for example? Is the edge of the occluding figure perceived as displaced such that it is seen as correspondingly shifted away from the partly occluded figure? Is the opposite edge of the partly occluded figure perceived as displaced outward to make room for the extra strip? Although these alternatives are logically possible, we suspect that the visual system somehow manages to see the partly occluded object as spatially extended perpendicular to the occluding edge without perceiving any difference in the positions of the regions attached to the edge. As strange as this may sound, the visual system manages to achieve equally bizarre results in other illusions. One example is the paradoxical motion one sees in motion aftereffects: Objects appear to move locally in a particular direction, but without changing their global positions. Perhaps the perceived extra strip due to partial modal completion is 'paradoxical' in much the same sense: It is there even though there is no well-defined position at which it is located.

References

Delboeuf J L R, 1892 "Sur une nouvelle illusion d'optique" *Academie Royale des Sciences, des Lettres et des Beaux Arts de Belgique. Bulletins* **24** 545–558

Helmholtz H von, 1867/1962 *Treatise on Physiological Optics* volume 3 (New York: Dover, 1962); English translation by J P C Southall for the Optical Society of America (1925) from the 3rd German edition of *Handbuch der physiologischen Optik* (first published in 1867, Leipzig: Voss)

Hershenson M (Ed.), 1989 *The Moon Illusion* (Hillsdale, NJ: Lawrence Erlbaum Associates)

Kanizsa G, 1979 *Organization in Vision: Essays on Gestalt Perception* (New York: Praeger)

Kanizsa G, Luccio R, 1978 *Espansione Fenomenica di Superfici in Condizioni di Completamento Amodale* Reports of the Institute of Psychology, University of Trieste, Trieste, Italy

Kaufman L, Rock I, 1962 "The moon illusion: I" *Science* **136** 953–961

Levitt H, 1971 "Transformed up–down methods in psycho-acoustics" *Journal of the Acoustical Society of America* **49** 467–477

Micali G, Giurissevich O, Serani M, 1978 *Il Completamento Amodale come Fattore dell'Espansione Fenomenica di Superficie* Reports of the Institute of Psychology, University of Trieste, Trieste, Italy

Nakayama K, Shimojo S, 1990 "da Vinci stereopsis: depth and subjective occluding contours from unpaired image points" *Vision Research* **30** 1811–1825

Palmer S E, 1999 *Vision Science: Photons to Phenomenology* (Cambridge, MA: MIT Press)

Perussia F, 1983 "L'espansione fenomenica in funzione del completamento amodale con superfici anomale" *Ricerche di Psicologia* **7** 187–193

Rock I, 1995 *Perception* (New York: Scientific American Library)

Rock I, Kaufman L, 1962 "The moon illusion: II" *Science* **136** 1023–1031

Vezzani S, 1999 "Shrinkage and expansion by amodal completion: a critical review" *Perception* **28** 935–947