

Stereoscopic Depth and the Occlusion Illusion

by

Stephen E. Palmer and Karen B. Schloss

Psychology Department, University of California, Berkeley

Running Head: Stereoscopic Occlusion Illusion

Send proofs to:

Stephen E. Palmer

Psychology Department

University of California, Berkeley

Berkeley, CA 94720-1650

(510) 642-7135

palmer@cogsci.berkeley.edu

Abstract

In the occlusion illusion, the visible portion of a partly occluded object appears larger than a physically identical non-occluded region. Stereoscopic displays allowed for a direct test of the apparent distance hypothesis. Experiment 1 measured both the perceived size and perceived depth of partly-occluded targets when the binocular disparity of both targets and occluders were varied. Stereoscopic occlusion greatly increased perceived target size but not perceived target distance. A reduced illusion was still present when the target was stereoscopically in front of the abutting rectangle, however. Experiment 2 gave similar results, even when the occluding figures were illusory rectangles that formed no explicit T-junctions. Experiment 3 showed that an unexpected negative size illusion on control trials was primarily due to adaptation to the occlusion illusion on other trials. The present findings eliminate apparent distance explanations of the occlusion illusion, but are consistent with other hypotheses, such as partial modal completion and selective dimensional expansion.

Kanizsa and Luccio (1978; translated in Kanizsa, 1979) first reported an illusion in which a figure that is perceived as partly occluded (the “target”) appears larger than a physically identical region that is fully visible (the “standard”). This illusion was later christened the “occlusion illusion” because of the key role played by evidence of occlusion (Palmer, 1999; Palmer, Brooks & Lai, 2007). Figure 1A shows a typical example in which a half-circle target abuts a rectangle along one edge, such that the display is perceived as a rectangle partly occluding a full circle. Under these circumstances, the partial-circle appears to be substantially larger than the optically identical standard shown next to it.




Figure 1. The occlusion illusion in flat and stereoscopic viewing conditions. (A) The partial-circle target that abuts the rectangle appears larger than the physically identical standard shown next to it. (B) When adjacent configurations are fused with crossed disparity, the rectangle in the left fused display appears stereoscopically closer than the target and the rectangle in the right fused display appears stereoscopically farther than the target. When adjacent configurations are fused with uncrossed disparity, the rectangle in the left fused display appears stereoscopically farther than the target and the rectangle in the right fused display appears stereoscopically closer than the target.

Palmer, Brooks, and Lai (2007) investigated this illusion in an attempt to determine the conditions under which it occurs and to test the predictions of two hypotheses about its cause. They found that the size of the illusion, which was sometimes as large as 20% in area, varied with the amount of evidence of occlusion, in that stronger evidence of occlusion produced larger illusions. They also examined two possible explanations – the partial modal completion hypothesis and the apparent distance hypothesis – and found stronger support for the former.

Palmer et al. (2007) argued that the occlusion illusion is of great potential interest because partial modal completion is a new mechanism underlying size illusions. Partial modal completion postulates that the illusion occurs because the visual system “fills in” a small strip of the partly occluded target surface along the occluding edge. The key idea is that when an object is perceived as occluded, not only is there the well-known and much-studied phenomenon of “amodal completion” of the target behind the occluder (i.e., perceptual completion without a corresponding sensory experience in the relevant modality; see Kanizsa, 1979), but there is also a small, but measurable, extension of sensory experience along the occluding edge for the partly occluded surface. Notice that this explanation implies that the occlusion illusion is actually an illusion of both size and shape because the target is perceived as though it were less occluded than it actually is.

There is an even more obvious explanation of the occlusion illusion in terms of apparent distance, however, that we attempt to rule out in the present experiments. The apparent distance hypothesis explains the occlusion illusion as yet another size illusion caused by applying the size/distance relation with an erroneous perception of distance, akin to the famous moon illusion (Rock & Kaufman, 1962; Kaufman & Rock, 1962). The logic of this explanation is as follows: When the target is perceived as occluded by the rectangle, it is necessarily at least somewhat farther away than the occluding rectangle. If the occluding rectangle is perceived as located in the same depth plane as the standard, then the target must also be at least somewhat farther away than the standard. According to the size/distance relation, the target should be perceived as correspondingly larger than the standard, because it subtends the same retinal angle yet is perceived as farther away. Note that the apparent distance explanation assumes that the illusion is purely one of size: as the target appears to be farther away, its extent is increased in all directions equally. This prediction contrasts with that of the partial modal completion hypothesis, which implies changes in both size and shape. This distinction was the basis for Palmer et al.’s (2007) explicit

test between these two explanations. Their data supported the partial modal completion hypothesis but were not strong enough to be definitive. Indeed, it is quite possible that both hypotheses are valid to some degree.

The viability of the apparent distance explanation, however, depends on precisely how much farther away the target is perceived to be than the standard, and precisely how much larger it is perceived to be than the standard. The primary motivation for studying the occlusion illusion under stereoscopic conditions is that the disparity-defined distances among the target, occluder, and standard can be precisely controlled so that measurements of both the perceived size and perceived distance of the target relative to the standard can be obtained. If the target and standard are perceived to be located in the same binocular depth plane as specified by equal disparity, for example, the apparent distance hypothesis predicts that the illusion should disappear because there is no difference in perceived distance. If the target is indeed perceived to be farther away than the standard, however, we can determine whether the increment in perceived distance is sufficient to account for the magnitude of the increment in perceived size. We test such predictions in the experiments described below.

We note in passing that yet another possible explanation for the occlusion illusion can be formulated in terms of “selective dimensional expansion.” Here the conjecture is that, for some as-yet-unidentified reason that relates to perceived occlusion, the visual system may selectively overestimate the dimension of the target figure that is perpendicular to the occluding edge. Like the partial modal completion hypothesis, selective dimensional expansion implies that the occlusion illusion has both a size and a shape component, because the target would appear stretched along its axis perpendicular to the shared edge. Palmer et al. (2007) did not test this possibility, but an extension of their Experiment 4 that is designed to do so is currently in progress. The difference between the shape change implied by partial modal completion (the insertion of an additional strip along the occluding border in the direction perpendicular to it) and selective dimensional expansion (stretching the entire target in the dimension perpendicular to the occluding border) is subtle, but potentially discernable.

In the present article, we report results that provide decisive evidence against the apparent distance hypothesis. In essence, we show that when binocular disparity is manipulated, a robust size illusion is still present under stereoscopic viewing conditions in which we find no evidence of any misperception of distance. Along the way we report two unexpected results. One is that monocular T-

junctions elicit the occlusion illusion, albeit reduced in magnitude, even when the abutting rectangle is clearly perceived to be stereoscopically behind the target (see Figure 1B). We find this effect even when the T-junctions are not physically present, but are themselves illusory. The second is that a *negative* occlusion illusion occurs when no occluding rectangle is present on control trials. That is, the target actually appears *smaller* than a physically identical standard when only those two figures are presented to the observer, provided that an occluding rectangle causes the target to appear larger on other (non-control) trials.

Experiment 1: Perceived Size and Distance in the Occlusion Illusion

Perhaps the most obvious testable prediction of the apparent distance hypothesis is that the partly occluded target must be perceived as farther away than the standard by a distance sufficient to account for the increase in its perceived size. Although the target necessarily appears at least slightly farther from the viewer than the rectangle that partly occludes it (see Figure 1), it is unclear whether it is also seen as noticeably farther away than the standard. If the apparent distance hypothesis is a viable explanation of the occlusion illusion, however, the target must be perceived as measurably farther than the standard. We tested this prediction in Experiment 1 by independently manipulating the stereoscopic distance to the two parts of the occlusion configuration – the target and the occluding rectangle, while keeping the standard in the plane of the monitor – and measuring both the target's perceived size (Experiment 1a) and its perceived distance (Experiment 1b) relative to the standard. If apparent distance explanations of the occlusion illusion are viable, then the target should be reported as farther away than its binocular disparity specifies by an amount that is sufficient to account for the increase in its perceived size.

Methods

Subjects. All 26 participants in Experiment 1a and all 11 participants in Experiment 1b were University of California, Berkeley undergraduates (mean age of about 20 years), who consented to participate for partial course credit in their psychology course. Data were excluded from five participants who did not pass the stereoscopic depth pre-screening test (see below). All gave informed consent, and the Committee for the Protection of Human Subjects at the University of California, Berkeley, approved the experimental protocol.

Design. There were 12 conditions defined by the orthogonal combination of occluder depth plane (near, middle, far, or absent) and target depth plane (near, middle or far).

Displays. Participants viewed the computer screen from approximately 70 cm. The monitor was 18" diagonally with a resolution of 800 x 600 pixels at 85 Hz. Objects were presented on a neutral gray background at 50% of maximum luminance. The occluder, target, and standard were pixilated with monochromatic noise, (mean contrast of 20%, 80% and 80%, respectively), to facilitate fusion of stereoscopic displays. The noise pattern for the target and standard differed so that their sizes could not be compared based on scrutiny of the fine-grained structure of their textures.

The target configuration (including the occluder, when present) and the standard were placed in a diagonal arrangement so that neither the horizontal nor vertical dimensions were aligned. The target configuration was located at the top-left of the screen and the standard at the bottom-right. The occluder and target were centered 50 px (1.87°) left of the vertical midline. The occluder was centered 20 px (0.75°) above the horizontal midline. The target was centered horizontally and placed directly above the occluder (202×112 px or $7.53^\circ \times 4.20^\circ$) with the target's 144 px diameter (5.37°) adjacent to occluder's top edge. The standard was located 50 px (1.87°) to the right of the vertical midline and 132 px (4.95°) below the horizontal midline. The center of the entire display was located at the center of the monitor. On each trial the target was stereoscopically positioned in one of three depth planes relative to that of the standard: near, same or far (total disparities of -16 px (0.60°), 0 px (0°), +16 px (0.60°), respectively). The occluder was either placed in one of the same three depth planes or was absent from the display.

The size of the target was increased/decreased during the adaptive staircase procedure by occluding/disoccluding rows of pixels behind the occluding rectangle rather than increasing/decreasing the target's radius, because prior research indicates that this transformation is a closer approximation to the distortion caused by the illusion (Palmer et al., 2007, Experiment 4). To be sure that no gap would appear between the target and the rectangle when disoccluding rows of it, an additional rectangular strip (144×48 px; $5.37^\circ \times 1.80^\circ$) was added below its straight edge, forming a bullet-shaped figure, the lower portion of which was occluded by the edge of the rectangle.

Procedure. Participants were pre-screened to determine whether they perceived depth accurately through the stereoscopic glasses by presenting them with three vertically aligned rectangles (the same rectangle used as the occluder during the experiment) and asking them to report the rectangle that

appeared closest and the one that appeared farthest. The top rectangle was in the *middle* depth plane, the center rectangle was *near* and the bottom rectangle was *far*. The results from five participants who reported the incorrect depth ordering were eliminated.

In Experiment 1a, a 1-up-1-down staircase procedure (Levitt, 1971) determined the point of subjective equality (PSE) in size between the target and standard. There was a separate staircase for each of the 12 conditions with randomly interleaved trials. Each display remained on the screen until the participant responded. Between trials, there was a 235ms inter-trial interval (ITI) during which a fixation cross was displayed in the center of the display. The target was smaller than the standard in the initial display for half of the staircases and larger for the other half, with the starting size being determined randomly.

Participants were instructed to indicate whether the target appeared larger or smaller than the standard by pressing one of two buttons on the mouse. If a participant indicated that the target appeared larger, then the next trial of that condition presented a smaller target by occluding one or more rows of pixels along its bottom edge; if he or she indicated that it appeared smaller, the next trial of that condition presented a larger target by adding one or more rows of pixels along its bottom edge.¹ The number of pixel-rows concealed or revealed started at 5 (0.19°) and decreased in increments of 1 (0.04°) over the course of reversals so that it was always a single pixel-row when each staircase converged. Each staircase was completed when participants made eight reversals. The PSE for each condition was computed as the mean value from the last six reversals of the staircase. The experimental program and images were adapted from those of Palmer, Brooks, and Lai (2007), which is archived in the *Journal of Neurobehavioral Experiments* and can be downloaded from http://nbs.neuroexpt.com/ex_files/expt_view?id=157.

In Experiment 1b, participants were asked to rate the distance to the target relative to the standard on a five-point scale ranging from -2 to +2 (-2 = far behind, 0 = same depth plane, +2 = far in front). There were three repetitions for each of the 12 stereoscopic configurations using targets that were physically identical to the standards, and another three repetitions for each of the 12 stereoscopic configurations using targets that were the same perceived size as the standards, as determined by the results of Experiment 1a. There were thus 72 trials in Experiment 1b.

Apparatus. The test items were generated in Adobe Photoshop and Illustrator. The procedure was programmed and presented in Presentation (www.neurobs.com). The stereoscopic glasses, model CE-WS E-2 Emitter, were manufactured by StereoGraphics® (www.stereographics.com).

Results and Discussion

Experiment 1a: Judgments of Perceived Size. Each staircase for each subject in Experiment 1a provided one estimate of the PSE for the size of the target relative to the size of the standard in the given experimental condition. The target PSE was calculated relative to the objective size of the standard in terms of the decrease in the number of pixels in the target's vertical extent needed to match the standard. This measure thus maps increases in the target's perceived size relative to the standard to higher values, consistent with the direction of the illusion (i.e., higher values correspond to cases in which the partly occluded target was perceived as larger than the standard). These measures were averaged over subjects and are plotted in Figure 2.

The size-distance relation implies that the no-rectangle conditions should lie on a single, ascending linear function that specifies the increase in perceived target size as a function of increasing binocular disparity. This effect is indeed present and approximately linear ($F(1,20) = 68.88$ $p < .001$), showing a clear influence of apparent distance of the target on its perceived size, independent of the perception of occlusion. The existence of this effect shows that observers' judgments of target size are indeed sensitive to size differences arising from differences in apparent distance, as required by the apparent distance hypothesis. Surprisingly, when the target is in the same (zero-disparity) plane as the standard, the perceived size of the unoccluded target is about 6 pixel-rows (0.23°) smaller than that of an identical standard. The same negative size effect appears to affect the perception of target size at the near and far distances as well, because the no-rectangle function is so nearly linear. We considered various hypotheses about why it might have arisen from various factors involving viewing conditions, but later discovered (see Experiment 3) that it is actually a context effect due to adaptation in the presence of other trials in which the occlusion illusion made the target appear larger when it was partly occluded by the rectangle. For now, we will simply call this seemingly anomalous finding the "negative bias effect" and largely ignore it until Experiment 3 addresses it directly.

In addition to binocular depth effects on the perceived size of targets, there are also clear and systematic effects due to occluder conditions ($F(3,60) = 38.53$, $p < .001$) and a sizable interaction between

target and occluder conditions ($F(6,120) = 14.75, p < .001$). These effects show that factors other than target disparity also influence the perceived size of the target. The inset in Figure 2 shows the results that would be predicted by a theory that posits effects due to both stereoscopic depth (arising from the binocular disparities of the target relative to the standard) and stereoscopic occlusion (arising from the binocular disparities of the target relative to the occluder), but nothing else. The stereoscopic occlusion effect should occur whenever the rectangle is stereoscopically in front of or coplanar with the target, because these are the conditions in which the rectangle's depth is consistent with it occluding the target. This (upward) increment in the rectangle conditions relative to the no-rectangle conditions should occur whenever the rectangle is as close as or closer than the target: i.e., for all three cases in which the rectangle is in the near plane, for the two cases in which the rectangle is in the middle plane when the target is in the middle or far plane, and for the single case in which the rectangle is in the far plane when the target is also in the far plane. This pattern is predicted by the apparent distance hypothesis, provided that only stereoscopic depth effects are present.




Figure 2. Results of Experiment 1a showing the perceived size of the target relative to the standard. Measurements of the PSEs for the size of the target relative to the standard are plotted

as a function of the target's stereoscopic depth when the occluder is located in the near plane (black squares), middle plane (dark gray triangles), far plane (light gray diamonds), and when the occluder is absent (white circles). The inset in the top right corner shows the results predicted by the apparent distance hypothesis if only on joint effects of stereoscopic depth and stereoscopic occlusion are present.

The observed data, plotted in Figure 2, generally conform to this pattern, except that all of the conditions containing a rectangle are displaced upward relative to the absent-rectangle conditions by a roughly constant amount, which we will call simply the "rectangle" effect. We modeled the results using multiple linear regression analysis, predicting the size data from target disparity (-1, 0, +1 for the near, middle, and far depth planes), stereoscopic occlusion (1 when the rectangle was in front of or in the same plane as the target and 0 when it was behind), and presence of the rectangle in the display (1 when the rectangle was present and 0 when it was absent). This model fit the data very well, accounting for 97% of the variance. Because the to-be-predicted variable is the number of pixel-rows occluded / disoccluded in the matched target, and the predictor variables are unitless, the beta weights for the predictors estimate the size of these effects in terms of the equivalent number of additional pixel-rows occluded in the target to produce size equivalence. The best fitting estimates are -12.7 pixel-rows for the close target disparity effect and +12.7 pixel-rows for the far target disparity effect, +8.6 pixel-rows for the stereoscopic occlusion effect, +5.4 pixel-rows for the rectangle effect, and -3.5 pixel-rows for the additive constant, which estimates the anomalous negative bias effect.

At least two facts are noteworthy in this pattern of results. First, relative binocular disparity indicating that the rectangle's depth is consistent with it occluding part of the target, produces a size illusion that was offset by occluding about 8.6 pixel-rows of the target. This is a sizable fraction (68%) of the size increase (12.7 pixel-rows) produced by the binocular disparity step between adjacent depth planes in this experiment. This implies that if the apparent distance hypothesis is to account for the increase in the target's perceived size due to stereoscopic occlusion, the target should be perceived to be about 68% of that depth increment farther than the standard when it is seen as occluded by the rectangle. We examine this prediction directly in Experiment 1b below.

Second, some additional monocular factor arising from the mere presence / absence of the rectangle also makes the target look larger, and it is offset by occluding about 5.4 pixel-rows of the target.

There are at least two different explanations of this monocular rectangle effect. One is that it is due to the mere presence of the rectangle in close proximity to the target, independent of occlusion, such that the apparent size of the target is increased by size assimilation due to the rectangle. Palmer et al. (2007) rejected an assimilation explanation of the occlusion illusion, because it was inconsistent with the effects of different sized rectangles. A more plausible account is that the rectangle effect arises from T-junctions that occur where the target abuts the rectangle independent of their binocular disparities. If this is indeed the relevant consideration, the visual system must register these monocular T-junctions as evidence that the target is occluded, even when it is stereoscopically (and perceptually) in front of the rectangle. This means that the presence of T-junctions increases the size of the illusion even though stereoscopic information conflicts with the perception of occlusion. This interpretation is consistent with Palmer et al.'s (2007) conclusion that stronger evidence of occlusion increases the magnitude of the occlusion illusion. We note in passing that the systematic increase in perceived target size that is present in the data as a function of increasing binocular depth of the target implies that participants saw the stereoscopic depth accurately and, therefore, that they did, in fact, see the target as unoccluded when it was in front of the rectangle. There is more direct evidence of this fact from the explicit distance judgments (see below). Nevertheless, monocular T-junctions may register outside awareness to increase the evidence for target occlusion whenever a rectangle abuts the target and thereby increase the size of the illusion. Experiment 2 will address this issue further by studying displays with illusory occluders, which produce no optical T-junctions with the target, but do produce perceptual T-junctions, and control displays that contain similar inducing elements but lack any sort of T-junctions.

Experiment 1b: Judgments of Perceived Distance. Versions of both the apparent distance hypothesis and the partial modal completion hypothesis are able to account for the pattern of size judgments in Experiment 1a, but they make very different predictions about corresponding judgments of perceived distance in Experiment 1b. The apparent distance hypothesis requires that the target be seen as farther away whenever it is partly occluded by the rectangle by an amount commensurate with its increase in perceived size. The partial modal completion hypothesis does not. What do the data reveal?

The average distance ratings reported by observers are plotted in Figure 3 as a function of target disparity, averaged over type of equivalence between target and standard (i.e., physical size and perceived size). As is evident by inspection there is a very large, highly linear effect of stereoscopic target

depth ($F(2,20) = 410.35$, $p < .001$). Neither type of size-equivalence nor occluder condition had a significant main effect ($F(1,10) = 1.25$, $p > .05$; $F(3,30) < 1$, respectively), but there was marginally reliable interaction among all three factors due to small variations in the near and far target depth planes ($F(6,60) = 2.69$, $p < .05$), the interpretation of which is unclear. In the crucial case when the target was in the middle depth plane, however, there were no differences of any kind in the perceived distance of the target relative to the standard. These findings seem to imply that the apparent depth hypothesis cannot account for the occlusion illusion because its explanation is based on observers perceiving the target as farther away when it is partly occluded by the rectangle than when the rectangle is absent. There is no hint of such effects in the data for the middle (zero disparity) depth plane.

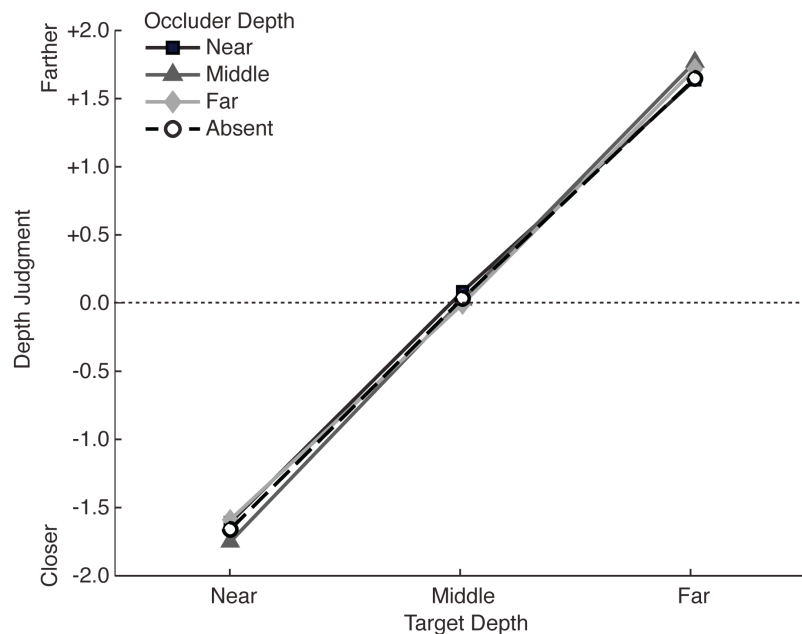


Figure 3. Results of Experiment 1b showing the perceived depth of the target relative to the standard. Target depth judgments are plotted as a function of target disparity when the occluder is located in the near plane (black squares), middle plane (dark gray triangles), far plane (light gray diamonds), and when the occluder is absent (white circles).

One possible objection to this conclusion would be that the amount of additional depth produced by occlusion is simply too small to be measured effectively in this experiment. The results of Experiment 1a, however, enable us to estimate the additional depth that would be required to induce the measured size illusion, and it is sizable. Using the results of the regression model of the data from Experiment 1a, we estimate that the stereoscopic occlusion illusion is equivalent to occluding approximately 14

additional pixel-rows from the physically identical target: 8.6 pixel-rows for the stereoscopic occlusion effect plus 5.4 pixel-rows for the monocular rectangle (T-junction) effect. The regression model also provides an estimate of the size difference that results from placing the unoccluded target in the near and far depth planes relative to the middle plane: 12.7 pixel-rows. This implies that for the apparent distance hypothesis to account for the magnitude of the occlusion illusion when the rectangle partly occludes the target, observers should perceive it to be slightly farther away than the unoccluded target in the next-farther depth plane. We also know from the distance judgment data that the observers perceived this depth difference very clearly. The depth increment between the near and middle planes in Experiment 1b was about 1.6 rating units out of 2.0, as was the depth decrement between the middle and far planes, and both of these disparity-based differences in depth ratings were extremely reliable ($F(1,10) = 235.92, 457.71$ $p = .0001$, respectively). In contrast, the purely occlusion-based difference in depth ratings did not come close to reaching statistical significance ($F(1,10) \leq 1$ for the target-near/occluder-near vs target-near/occluder-absent comparison, the target-mid/occluder-mid vs. target-mid/occluder-absent comparison, and the target-far/occluder-far vs. target-far/occluder-absent comparisons). Even if the rectangle effect noted in Experiment 1a were actually due to size assimilation rather than occlusion, the relative stereoscopic occlusion effect is still 8.6 pixel-rows, which is two-thirds of the size increase induced by the disparity step between depth planes. It thus seems highly unlikely that the present depth judgments were not sensitive enough to detect such large predicted differences in perceived depth, if they were indeed present.

Another possible objection to the conclusion that the present results rule out explanations in terms of apparent distance is suggested by Kaufman and Rock's (1948; Rock & Kaufman, 1948) two-stage theory of why the moon appears closer when it is low in the sky at the same time as it appears to be larger. They argue that the moon's larger perceived size in a first stage of non-conscious distance perception (where the size illusion occurs) makes it appear to be closer in a second stage of conscious distance perception (where its size makes it seem closer). In the present case, the distance increments due to occlusion and T-junctions, which are predicted by the apparent distance hypothesis, might fail to be perceived because the larger perceived size of the target makes it appear to be closer than it otherwise would. In order for this account to predict the obtained finding of zero depth effects, the increase in the perceived size of the target due to the occlusion illusion would have to exactly offset the increment in

perceived distance predicted by the apparent distance hypothesis. This argument is undercut in the present results, however, by the fact that the perceived depth data were essentially identical when the targets were psychophysically matched to the standard so that they were perceived as being the same in size. In these size-matched conditions, observers saw the sizes of the target and standard as being essentially the same, so differences in perceived size could not influence their judgments of perceived depth. Nevertheless, we will return to these issues in Experiment 2, which provides an even more stringent test of whether partial occlusion induces any increment in the perceived distance of the target relative to the standard.

Experiment 2: Effects of Real versus Illusory Occluders

The second experiment was undertaken to explore two findings from Experiment 1 in more detail. In Experiment 2a we examined the relative effects of explicit (real) versus implicit (illusory) T-junctions on the rectangle effect of the size illusion measured in Experiment 1a. In Experiment 2b we made more sensitive measurements of perceived distance of the target relative to the standard that were designed to reveal even very small depth differences arising from perceived occlusion that might have been missed in Experiment 1b.

The role of T-junctions was explored by including two new occluder conditions in addition to the real-rectangle and no-rectangle conditions used in Experiment 1: an illusory rectangle condition, in which the occluder was induced by edge alignment of other partly occluded elements, and a control for this display, in which the inducing elements were misaligned to block perception of the illusory rectangle (see Figure 4). In the illusory rectangle condition, the edges of the occluding rectangle were not physically present, but merely induced by contextual figures that also appeared to be occluded by a rectangle that was the same color as the background (Kanizsa, 1979). No explicit T-junctions were present in this condition. Both Perussia (1983, as described in Vezzani, 1999) and Palmer et al. (2007) reported that such occluders do, in fact, produce an occlusion illusion, albeit a smaller one than is produced by an equal-sized "real" rectangle defined by luminance differences. The question of primary interest was whether such illusory occluders would continue to produce an occlusion illusion, like the real occluder did, when binocular disparity specifies that the inducers of the illusory rectangle are farther away than the target. If explicit T-junctions are required, then the illusory occluder should produce no occlusion illusion in such

depth conditions. If implicit T-junctions are sufficient, then we expect a reduced occlusion illusion to be present.




Figure 4. Sample displays for Experiments 2a and 2b. The four occluder conditions in Experiment 2 were: (A) real rectangles, (B) no-rectangle controls, (C) illusory rectangle controls, and (D) illusory rectangle controls.

Perceived depth was measured in a separate procedure (Experiment 2b) in which the binocular disparity of the target was varied over a small range near the depth plane of the monitor, at which the standard was always positioned. On each trial the target was stereoscopically placed at one of five slightly different depths relative to that of the standard (total disparities of -4 px (0.15°), -2 px (0.07°), 0 px (0°), 2 px (0.07°), and 4 px (0.15°)), all of which were much less than the binocular depth of the occluding rectangles in Experiment 1 (total disparities of -16 px (0.60°), 0 px (0°), and $+16$ px (0.60°)). Observers rated the depth of each target relative to the standard using a nine-point scale. These disparities were small enough to ensure detection of even minor depth differences due to occlusion, if such occurred.

Methods

Participants. All 12 participants (mean age of about 20) were University of California, Berkeley, undergraduates who consented to participate for partial course credit in their psychology course.

The same observers participated in Experiments 2a followed by 2b.

Experiment 2a: Size Judgment

Design. There were 20 conditions defined by the orthogonal combination of the following factors: occluder depth plane relative to the monitor (near, middle, or far), occluder type (real rectangle, illusory rectangle, illusory rectangle control), and target starting size (smaller or larger than the standard). A no-rectangle control was also included with both the smaller and larger starting sizes. Half of the participants viewed the configural display at the top of the screen and the standard figure on the bottom, and the other half viewed them in the reverse positions. The target always appeared at zero disparity in the depth plane of the monitor and standard.

Displays. The displays were similar to those in Experiment 1 except for the three different types of occluders (see Figure 4). The “real” rectangular occluder was solid dark gray (20% of maximum luminance) with a texture of single-pixel white dots. The inducing shapes that formed the illusory-rectangle occluder and the illusory rectangle control were also solid dark gray (20%). The shapes in the illusory rectangle control were the same as in the illusory rectangle composition except that they were rotated or displaced to eliminate the alignment of adjacent edges that induce the perception of illusory contours. The two partial-circles were light gray (80% of maximum luminance) with a texture of single-pixel black dots. The pattern of the textural dots was different on the two partial-circles to eliminate strategies based on matching specific texture elements.

The real rectangle and the illusory rectangle were both 202 x 112 px (7.53° x 4.20°) in size, and both were the same size as the occluders in Experiment 1. The standard partial-circle was also the same size as in the previous experiments, and the target size varied in the same way, with the same extra strip forming the “bullet” shape. The entire configurations that formed the illusory rectangle and illusory rectangle control were both (324 x 124 px (12.06° x 4.65°) in size.

Half of the participants viewed the occlusion configuration in the top left portion of the monitor, which was in the same position as in Experiment 1. The other half viewed them in the reverse positions, in which the configural display was in the bottom right portion of the monitor (occluder and target centered 50 px (1.87°) right of the vertical midline, occluder 204 px (7.64°) below the horizontal midline and target varying horizontally during the adaptive procedure) and the standard was in the top left portion (centered 50 px (1.87°) to the left of the vertical midline and 112 px (4.20°) above the horizontal midline). In both cases the center of the entire display was located at the center of the monitor and the

diagonal arrangement described in Experiment 1 was maintained. Both partial circles were coplanar with the monitor in every condition; only the occluder varied in depth.

Procedure. The procedure for Experiment 2a employed the same staircase algorithm as in Experiment 1a. Testing began with the preliminary depth test, which was followed by 20 interleaved staircases.

Experiment 2b: Depth Judgments

Design. The depth judgment portion included the orthogonal combination of the 20 display conditions described above with five different target depth positions relative to the monitor (-4, -2, 0, +2, and +4 pixels of disparity, where negative numbers indicate positions behind the monitor plane and positive numbers indicate positions in front of the monitor plane). Participants viewed the objects in the depth judgment part of experiment in the same arrangement on the monitor as they had in the size judgment part of the experiment.

Displays. The displays were the same as Experiment 2a except that the target was always the same physical size as the standard (since this did not matter in Experiment 1b) and could occupy any of five positions in depth: -4, -2, 0, +2, and +4 pixels of total disparity. The rectangle also varied in depth as it did in Experiment 2a.

Procedure. Participants rated the target's depth relative to the standard on a nine-point scale (-4 = far behind, 0 = same depth, +4 = far in front). Trials were presented in random order with two repetitions of each.

Result and Discussion

Experiment 2a: Size Judgments. PSEs for target sizes relative to standard sizes were computed as in Experiment 1. The results are shown in Figure 5, averaged over target starting sizes (smaller or larger than the standard) and configuration positions (configural display at the top left or the bottom right) because there neither factor showed an effect ($F(1,11) = 3.88, p > .05, F < 1$, respectively) or interacted reliably with any other factors. An overall analysis of variance showed main effects of occluder condition ($F(2,22) = 36.21, p < .001$), occluder disparity ($F(2,22) = 110.53, p < .001$), and an interaction between these two variables ($F(4,44) = 4.25, p < .01$).

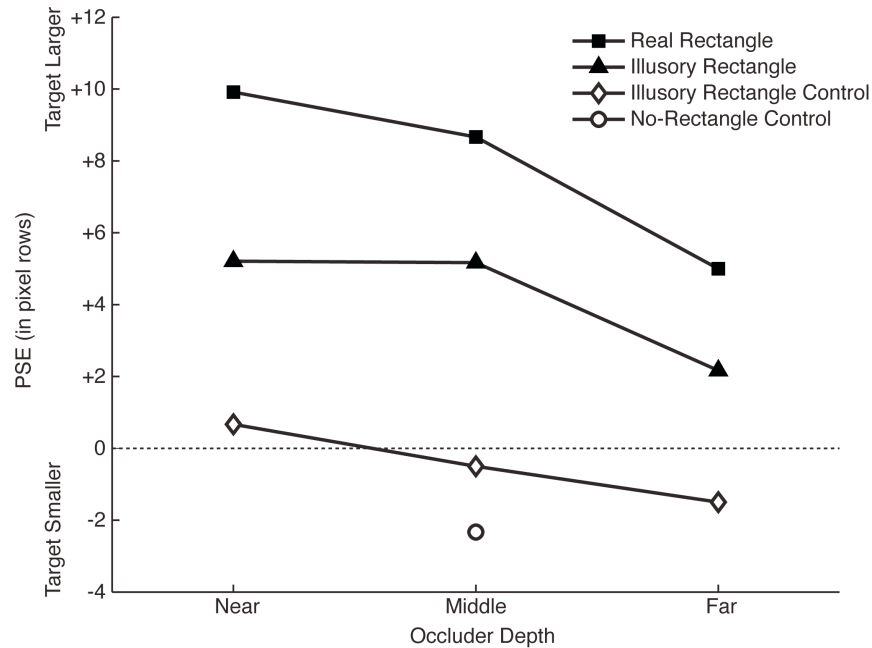


Figure 5. Results of Experiment 2a showing the perceived size of the target relative to the standard. Measurements of the PSEs between the target and standard are plotted as a function of occluder depth for the real rectangle condition (black squares), illusory rectangle condition (black triangles), illusory rectangle control condition (white diamonds), and no-rectangle control condition (white circle).

The results for the real rectangle occluder are consistent with those for the corresponding conditions in Experiment 1a: a large occlusion illusion of about 11 pixel-rows was present when the rectangle was stereoscopically in front of the target in the near plane ($F(1,11) = 67.90$, $p < .001$) and a similar sized illusion of about 9 pixel-rows was present when the rectangle was stereoscopically coplanar with the target in the middle plane ($F(1,11) = 54.42$, $p < .001$). These effects differed only slightly from each other ($F(1,11) = 5.16$, $p < .05$). A reduced but reliable, illusion of about 5 pixel-rows was evident when the rectangle was stereoscopically behind the target in the far plane ($F(1,11) = 33.86$, $p < .001$), compared to when the rectangle was in the near ($F(1,11) = 16.54$, $p < .01$) and middle ($F(1,11) = 10.22$, $p < .01$) depth planes. The latter replicates the monocular rectangle effect evident in the results of Experiment 1a.

The size illusion in the illusory occluder conditions was weaker than in the real occluder conditions ($F(1,11) = 20.47$, $p < .01$), but the pattern was similar. A significant illusion of about 5 pixel-rows (0.19°) was present when the illusory rectangle was stereoscopically in front of the target in the near

plane ($F(1,11) = 34.17, p < .001$) or coplanar with the target in the middle plane ($F(1,11) = 31.26, p < .001$), and these effects did not differ significantly from each other ($F < 1$). Nevertheless, a reduced, but still reliable, illusion was evident when the illusory rectangle was stereoscopically behind the target in the far plane (about 2 px, $F(1,11) = 12.94, p < .01$), compared to when the rectangle was in the near ($F(1,11) = 12.42, p < .01$) and middle ($F(1,11) = 11.37, p < .01$) depth planes.. In contrast, the illusory control conditions produced no reliable illusion relative to the theoretical zero point ($F < 1$). Nevertheless, the illusory control condition in the middle depth plane was about 2 pixel-rows greater than the no-occluder control condition ($F(1,11) = 5.68, p < .05$). As in Experiment 1a, there was a significant difference between the no-rectangle control condition and the theoretical zero point ($F(1,11) = 6.58, p < .05$). The fact that the present experiment was counterbalanced for the position of the standard and configural portion of the display means that this negative bias effect is not due to positional differences. We will return to the question of its cause in Experiment 3.

We modeled the results using multiple linear regression analysis as in Experiment 1a. The size data was predicted from stereoscopic occlusion (1 when the contextual element – real rectangle, illusory rectangle or illusory rectangle control – was in front of or in the same plane as the target and 0 when it was behind), presence of T-junctions (1 if real or illusory, 0 if absent), occluder visibility (1 if real, 0 if illusory or absent), and presence of contextual elements (1 if a real rectangle, illusory rectangle or illusory rectangle control was present, 0 if they were absent). The model fit the data extremely well, accounting for 98% of the variance. The best fitting estimates were 3.0 pixel-rows for the stereoscopic occlusion effect, 4.6 pixel-rows for the T-junction effect, 3.7 pixel-rows for occluder visibility, -0.09 pixel-rows for presence of contextual elements and -2.3 pixel-rows for the additive constant, which measures the negative size bias.

The results of Experiment 2a thus replicate the corresponding conditions in Experiment 1a and extend them by showing that an illusory rectangle, which contains no explicit T-junctions when abutting the target, produces effects similar to, but smaller than, those of a real rectangle with explicit T-junctions. This pattern of results suggests that T-junctions are an important factor in the occlusion illusion and that illusory T-junctions simply provide weaker evidence of occlusion than explicit T-junctions. Interestingly, both explicit and illusory T-junctions appear to produce a reliable occlusion illusion even when disparity information indicates that the abutting rectangle lies stereoscopically behind the target. We were

somewhat surprised by this result, because stereoscopic information can be important in determining the perception of implicit T-junctions at illusory contours, even to the point of blocking the phenomenal perception of those monocular implicit T-junctions as T-junctions in depth (Gillam & Nakayama, 2002). Nevertheless, there are monocular cues for illusory T-junctions at illusory contours in both the left-eye and the right-eye images of our displays taken singly, and these monocular cues appear to influence the perceived size of the target even when the target is phenomenally perceived as located in front of the illusory rectangle. This finding is consistent with previous results showing that the magnitude of the occlusion illusion is sensitive to the strength of the evidence for occlusion (Palmer et al., 2007), provided that we make the further reasonable assumptions that (a) implicit T-junctions are weaker cues to occlusion than explicit T-junctions and (b) monocular evidence of T-junctions is a weaker cue to occlusion than binocular evidence of T-junctions. Implicit monocular T-junctions (i.e., T-junctions formed by a monocular illusory contour and an abutting edge) thus appear to provide weak evidence of occlusion that is nevertheless sufficient to induce a measureable occlusion illusion, even when the phenomenal percept is that the monocularly abutting edge of the target lies in front of the illusory contour.

Experiment 2b. Recall that the target was located at five different distances relative to the standard, defined by binocular disparities of -4 px (0.15°), -2 px (0.07°), 0 px (0°), 2 px (0.07°), and 4 px (0.15°), a disparity range that was only 25% of that for the occluders (-16 px (0.60°), 0 px (0°), and +16 px (0.60°)). The results, plotted in Figure 6, show the presence of the expected monotonic increases in perceived distance as a function of target disparity ($F(4,44) = 52.56, p < .001$). There was no main effect of occluder condition ($F < 1$) or of occluder disparity ($F < 1$) and no interaction between occluder disparity and target disparity, ($F(6,66) = 1.60, p > .05$). There was a small but significant interaction, however, between target disparity and occluder condition ($F(12,132) = 2.49, p < .01$). Somewhat surprisingly, the real rectangles behaved in the same way as the no-rectangle control condition ($F < 1$) and the illusory rectangles behaved the same as the illusory rectangle controls ($F < 1$). Relative to the real rectangles and no-rectangle controls, the illusory rectangles and their controls both made the target appear slightly farther from the standard at non-zero disparities. This is illustrated in Figure 6 by the somewhat steeper slopes of the functions for the latter two conditions than for the former two. It is not obvious why this might be the case.

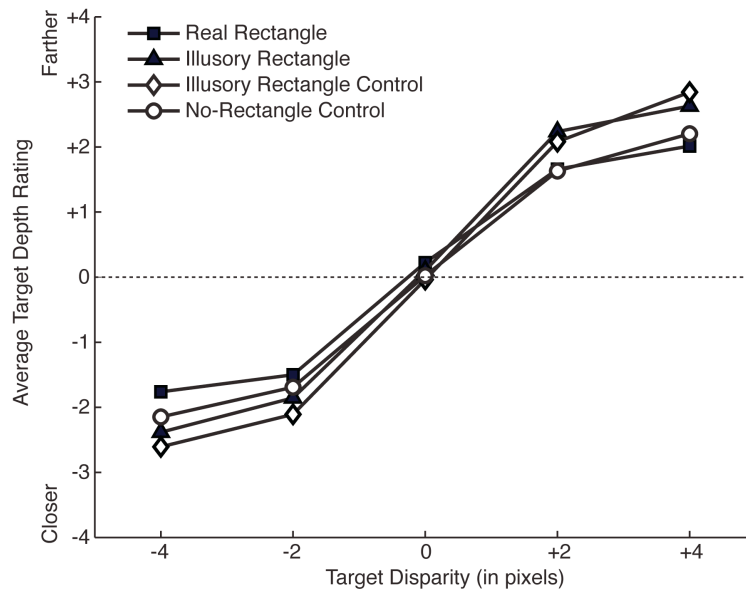


Figure 6. Results of Experiment 2b showing the perceived depth of the target relative to the standard. Target depth judgments are plotted as a function of target disparity for the real rectangle condition (black squares), illusory rectangle condition (black triangles), illusory rectangle control condition (white diamonds), and no-rectangle control condition (white circles).

The crucial question, however, is whether there are any measurable differences in perceived distance for the target when it is at the same disparity as the standard (0 pixels) due to different occluder conditions, because these are the conditions under which the occlusion illusion typically appears. To be even qualitatively consistent with the results of Experiment 2a for perceived size in this zero disparity condition in which the occluder is in the middle depth plane (see Figure 5), the apparent distance hypothesis requires that the perceived distance of the target be significantly farther than the standard for both the real rectangle and illusory rectangle conditions and that this depth effect be greater for the real rectangle occluder than for the illusory rectangle occluder. None of these effects are present in the data at zero disparity, however ($p > .05$, in all cases). The absence of such effects cannot be due to insensitive measures of depth, because we found highly reliable differences in perceived depth at disparities as small as ± 2 pixels (0.07°), ($F(1,11) = 35.48, 88.44 p < .0001$). The observers in Experiment 2 were thus able to detect a depth difference, with high reliability, that is roughly $1/6$ as large as the depth difference that the apparent distance hypothesis requires to account for the size of the occlusion illusion (using the data from Experiment 1). The illusion thus appears to be fully present when observers perceive the target to be

occluded, yet coplanar with the standard, as if overlapping pieces of paper were glued together. We therefore conclude that the apparent distance hypothesis of the occlusion illusion simply is not tenable.

Experiment 3: Inter-trial Contrast Effects on No-Rectangle Control Conditions

In both Experiments 1 and 2 there was a no-rectangle control condition in which only the two semicircles (the target and the standard) were on the screen and both were at zero disparity. The obvious expectation is that the PSE between target size and standard size in Experiments 1a and 2a would be the point at which they were physically equal (zero), because there was nothing else on the screen that could have influenced the target's perceived size. Nevertheless, in both experiments the results indicated that people saw the target as reliably *smaller* than the standard in this control condition – about 5 pixel-rows smaller in Experiment 1a and about 2 pixel-rows smaller in Experiment 2a – an effect in the direction opposite of the occlusion illusion. We found this unexpected and persistent negative bias effect puzzling and considered several possible explanations. We initially thought that it might somehow be due to the locations of the target configuration and standard on the monitor in Experiment 1a relative to the observer's eye, but when we varied their relative positions in Experiment 2a, we found no differences due to this factor.

Since there was nothing about the displays in the no-rectangle control trials themselves that could account for this odd effect, we speculated that it might be the result of an inter-trial context effect induced by target-size adaptation from trials that did include the rectangle. If so, there are at least two hypotheses about how the rectangle might cause the effect. One account is that the visual system adapts to some degree to the larger perceived size of the target in the occluded configurations due to the occlusion illusion itself, such that when the occluder is absent, this adaptation causes the target to be perceived as smaller than it actually is. Another hypothesis is that the adaptation is caused simply by the rectangle's presence adjacent to the target through a size assimilation effect, such that when the occluder is absent, the target looks smaller than it actually is due to the lack of the assimilation effect. The key difference between these two explanations is whether the rectangle needs to partly occlude the target for the negative contextual effect to occur or whether its mere presence in proximity to the target is sufficient.

In the present experiment we looked explicitly for an inter-trial context effect that influenced the no-rectangle control condition by manipulating whether the observers ever saw a rectangle or not. If they

did, we controlled whether it appeared to occlude the target or was merely adjacent to it. There were three primary conditions in the between-participant design we used: a no-rectangle condition in which no rectangle ever appeared on any trial, an occluding-rectangle condition in which half of the trials contained an occluding rectangle and half contained no rectangle, and an occluded-rectangle condition in which half of the trials contained a rectangle that was stereoscopically behind and partly occluded by the target, and half contained no rectangle.

If the negative size effect on no-rectangle control trials of Experiments 1a and 2a is an inter-trial context effect, then no such effect should be present in the no-rectangle condition of the present experiment because participants would never have experienced trials in which a rectangle of any sort was present. If the effect is an inter-trial adaptation effect due specifically to the occlusion illusion on other trials (i.e., arising from partial modal completion), then it should appear in the occluding-rectangle condition, where the target is occluded, but not in the occluded-rectangle condition, where the target occludes the rectangle. If it is due to adaptation resulting from size assimilation effects from a nearby rectangle independent of occlusion, then it should be equally present in both the occluding-rectangle and the occluded-rectangle conditions. If it is due to both occlusion and size assimilation, then it should be present in both rectangle conditions, but greater in the occluding-rectangle condition than in the occluded-rectangle condition.

Methods

Participants. All 36 participants were students at the University of California, Berkeley, who received partial course credit in their undergraduate psychology course. They were equally divided and randomly assigned among the three conditions (no-rectangle, occluding-rectangle, and occluded-rectangle). All gave informed consent, and the University of California, Berkeley Committee for the Protection of Human Subjects approved the experimental protocol. Their mean age was about 20 years.

Design. All three primary conditions included six staircases for the orthogonal combination of three target depths (near, middle, far) and two starting sizes (smaller or larger than the standard). In the rectangle conditions, another six analogous staircases contained the rectangle. In the occluding-rectangle condition the rectangle abutted the straight side of the target such that it appeared to partly occlude the target, as in previous experiments. In the occluded-rectangle conditions the rectangle was displaced upward toward the target such that the target overlapped it and appeared to partly occlude the rectangle.

Displays. The target and standard in all three conditions were the same objects located in approximately the same screen positions as in Experiment 1a. The target appeared in the near, middle or far depth planes (total disparity of -16 px (0.60°), 0 px (0°), +16 px(0.60°)). In the occluding-rectangle condition the rectangle, when present, was slightly in *front* of the target stereoscopically (total disparities of -18 px (0.67°), -2 px (0.07°), +14 px (0.52°)) and its top side abutted the bottom side of the semicircular target. In the occluded-rectangle condition, the rectangle was shifted vertically toward the target by 10 px (0.38°) so that the target overlapped it, and it was positioned slightly *behind* the target in depth (total disparity of -14 px (0.52°), +2 px (0.07°), +18 px (0.67°)) to decrease the chances that participants would perceive it as occluding the target.

Procedure. The procedure was the same adaptive staircase procedure as in Experiments 1a and 2a.

Results and Discussion

The primary results are plotted in Figure 7 as a function of target disparity for the conditions averaged over starting target sizes, which had no reliable effect in any condition. In the no-rectangle condition, there was only the expected main effect of perceived size as a function of target depth plane due to the size-distance relation ($F(2,22) = 7.21, p < .01$). Crucially, when the target was in the same depth plane as the standard (the middle, zero-disparity depth plane), the PSE in the no-rectangle condition was not significantly different from zero ($F(1,11) = 2.59, p > .05$). This fact implies that the negative PSEs measured in Experiments 1a and 2a were indeed due to some sort of contextual effect arising from the presence of the rectangles on other trials of the same experiment, because no such effect occurs in the no-rectangle condition of the present study.

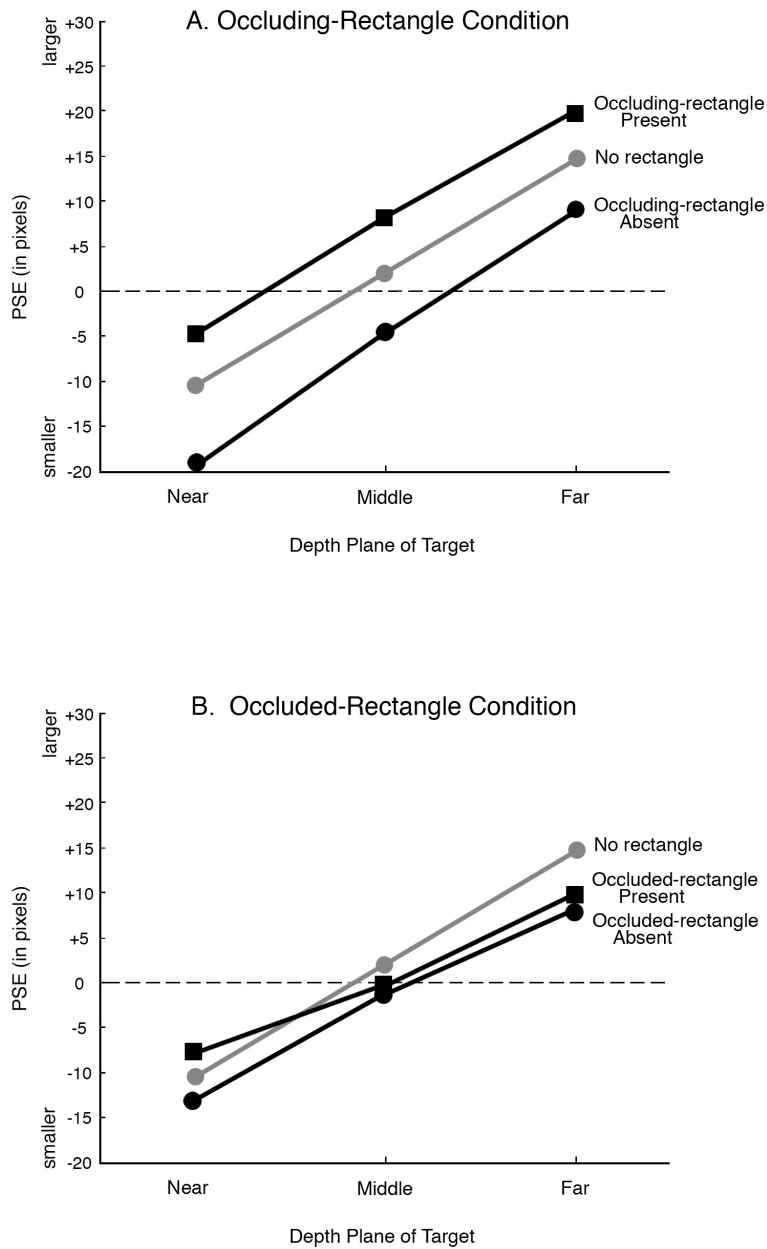


Figure 7. Results from Experiment 3 showing the perceived size of the target relative to the standard. (A) The occluding-rectangle configuration is shown at the top left of the graph. Measured PSEs for the target relative to the standard are plotted as a function of the binocular disparity of the target when the occluding-rectangle is present in that display (black squares), when the occluding rectangle is absent on that trials but present on other trials (black circles), and

when no rectangle is ever present in any display (gray circles). (B). The occluded-rectangle configuration is shown at the top left of the graph. Measured PSEs for the target relative to the standard are plotted as a function of the binocular disparity of the target when the occluded rectangle is present on that trial (black squares), when the occluded rectangle is absent on that trial but present on other trials (black circles), and when no rectangle is ever present in any display (gray circles).

The results of the occluding-rectangle condition were clear and straightforward. There was the same linear effect of target depth plane as in the no-rectangle condition ($F(2,22) = 37.15, p < .001$), and there was also a large main effect of rectangle presence/absence in the occluding rectangle condition ($F(1,11) = 98.89, p < .001$). The latter consisted of the combination of a positive size bias that made the target look bigger (i.e., an occlusion illusion) when the rectangle partly occluded it ($F(1,22) = 5.50, p < .05$, relative the corresponding no-rectangle conditions) and a negative size bias that made the target look smaller (i.e., an occlusion-illusion contrast effect) when the rectangle was not present ($F(1,22) = 10.78, p < .01$, relative to the corresponding no-rectangle conditions). Indeed, there was no interaction between target disparity and presence/absence of occluding rectangle in the occluding-rectangle condition ($F(2,22) = 3.17, p > .05$) or in these conditions relative to the no-rectangle control condition ($F < 1$). Thus, the results from the occluding-rectangle conditions suggest that an occluding rectangle produced an additive increment in the perceived size of the target on trials when the rectangle was present and an equally large decrement in the perceived size of the target when it was absent. This is what one would expect if the visual system adapts to the larger perceived size of targets when they are occluded, such that on trials in which the rectangle is absent, the target is perceived as smaller than it actually is.

The results were less straightforward in the occluded-rectangle conditions. The main effect of target depth plane was again present ($F(2,22) = 30.70, p < .001$). Although the slope of the function appears to be slightly shallower in this condition than in the no-rectangle condition, this interaction is not statistically reliable ($F < 1$). There was a small main effect of rectangle presence/absence, such that the target was seen as very slightly larger when the rectangle was present than when it was absent ($F(1,11) = 6.33, p < .05$), but there was no difference between the rectangle-present and no-rectangle conditions ($F < 1$) or the rectangle absent and no rectangle conditions ($F(1,22) = 4.21, p > .05$). Additionally, there was no interaction between the no-rectangle condition and rectangle present or absent ($F < 1$). Target depth plane

and rectangle presence/absence also interacted in the occluded-rectangle condition, such that the target looked larger when it occluded the rectangle than when the rectangle was absent if it was in the near and far planes ($F(1,11) = 9.36, 5.16, p < .05$), but not when it was in the same depth plane as the standard ($F < 1$). This slight size increase could result from the target being interpreted as a full circle that is partly inserted into a slit in the contextual rectangle, and thereby partly occluded (cf. Palmer, Brooks & Lai, 2007). Most of these effects are quite small, however, when compared to the corresponding effects in the occluding-rectangle conditions.

The PSE for the neutral configuration in each condition (the case in which the occluder was absent and the two partial circles were coplanar in the zero disparity plane) was compared with zero to test specifically for inter-trial context effects. In the no-rectangle condition it was not significantly different from zero ($F(1,11) = 2.59, p > .05$). The neutral conditions were reliably less than zero in both the occluding-rectangle and the occluded-rectangle conditions ($F(1,11) = 36.78, 6.65, p < .001, .05$). The fact that both of the conditions in which rectangles were present on other trials were different from zero suggests that the rectangle's presence on other trials includes a slight size assimilation component. The inter-trial contextual effect was significantly larger in the occluding-rectangle condition than in the occluded-rectangle condition ($F(1,22) = 7.67, p < .05$), however, indicating that there is an additional size contrast component due to the rectangle occluding the target.

Summary and Conclusions

The present results provide powerful constraints on possible explanations of the occlusion illusion. First, they show that the magnitude of the illusion under stereoscopic conditions depends on the relative binocular depth of the occluder and the target: The occlusion illusion is significantly larger when the occluder is in front of or coplanar with the target than when it is behind the target. This implies that the perception of the target being partly occluded by the rectangle is indeed important for the magnitude of the illusion. Second, they show that even under stereoscopic conditions, the magnitude of the illusion increases with the strength of the evidence for occlusion. Evidence of occlusion includes not only the relative binocular disparity of the target and occluder, but explicit monocular T-junctions (with real occluders) and, to a lesser extent, implicit monocular T-junctions (with illusory occluders). Third, they show that an inter-trial contrast effect is present on no-occluder trials such that the target is perceived as

smaller than the standard even when no other element is present in the display. This context effect appears to be primarily driven by the larger perceived size of targets due to the presence of the occlusion illusion on trials that do contain an occluding figure. Fourth, and most importantly, the present results show that, when the apparent depth of the target relative to the standard is actually measured in occluded versus unoccluded conditions, the target appears no farther away when the occluder is present than when it is not; only the binocular disparity of the target influences the perceived depth of the target relative to the standard. This finding conclusively shows that apparent distance is not a viable explanation for the occlusion illusion. An alternative hypothesis that does not depend on differences in apparent distance, such as partial modal completion or selective dimensional expansion, will be required to explain this illusion.

References

- Gillam, B., & Nakayama, K. (2002). Subjective contours at line terminations depend on scene layout analysis, not image processing. *Journal of Experimental Psychology: Human Perception and Performance*, 28, 43–53.
- Kanizsa, G. (1979). *Organization in vision: Essays on Gestalt perception*. New York: Praeger.
- 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.
- 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.
- Palmer, S. E. (1999). *Vision Science: Photons to Phenomenology*. Cambridge MA: MIT Press.
- Palmer, S. E., Brooks, J. L., & Lai, K. S. (2007). The occlusion illusion: Partial modal completion or apparent distance? *Perception*, 36, 650-669.
- Perussia F. (1983). "L'espansione fenomenica in funzione del completamento amodale con superfici anomale." *Ricerche di Psicologia*, 7, 187-193.
- 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.

Author note

Correspondence should be addressed to Stephen E. Palmer, 3210 Tolman Hall,
University of California, Berkeley, CA 94720-1650. E-mail: palmer@cogsci.berkeley.edu.

We thank the members of PalmerLab for their help and support during the conduct of this research and Johannes Burge for his comments on a previous version of this manuscript. We also thank Jeremy Wolfe and two anonymous reviewers for their helpful comments on a previous version of this report.
