[*Journal of the Optical Society of America A*, 13:681–688, April 1996]

Occlusion edge blur: A cue to relative visual depth

Jonathan A. Marshall,* *Christina A. Burbeck*,*† *Dan Ariely*,†
Jannick P. Rolland,* and *Kevin E. Martin**
Departments of *Computer Science, CB 3175, and †Psychology, CB 3270,
University of North Carolina, Chapel Hill, NC 27599, U.S.A.

(Correspondence to: Dr. Jonathan A. Marshall
919-962-1887, fax 919-962-1799, marshall@cs.unc.edu)

Abstract

We studied whether the blur/sharpness of an occlusion boundary between a sharply focused surface and a blurred surface is used as a relative depth cue. Observers judged relative depth in pairs of images that differed only in the blurriness of the common boundary between two adjoining texture regions, one blurred and one sharply focused. Two experiments were conducted; in both, observers consistently used the blur of the boundary as a cue to relative depth. However, the strength of the cue, relative to other cues, varied across observers. The occlusion edge blur cue can resolve the near/far ambiguity inherent in depth-from-focus computations.

Introduction

Faced with the task of inferring the 3D geometry of a scene from 2D images, the human observer uses a wide range of cues. In single 2D images, occlusion, perspective, shading, shadows, texture gradients, and familiar size can all provide information about 3D geometry; stereo comparison of two 2D images provides another cue. Most of these cues have been known for centuries. Here we identify a 2D image feature which has not been included in lists of depth cues[1,2] but which can influence the perception of depth.

Fig. 1 illustrates the cue. The two images in this figure are identical except for the nature of the edge between the inner and outer squares in each image. In Fig. 1a, the edge is sharp; in Fig. 1b, it is blurred. In experiments described below, we study the effect of this difference, which we call *occlusion edge blur*, on the relative depths seen in the two images.

In an image of a 3D scene, the blurriness of a surface provides information about the distance in depth between the surface and the observer's point of accommodation, but it does not, by itself, provide information about the direction of that difference in depth.[3,4] A blurry surface could be nearer to or farther from the observer than a sharply focused one. In some conditions, however, the availability of other types of information about the probable arrangement of the scene can render blur useful as a cue to relative depth, at least in theory. In particular, when two regions share a common boundary, implying that one is likely to be partially occluding the other, the edge between them may carry information about which region is the occluder.

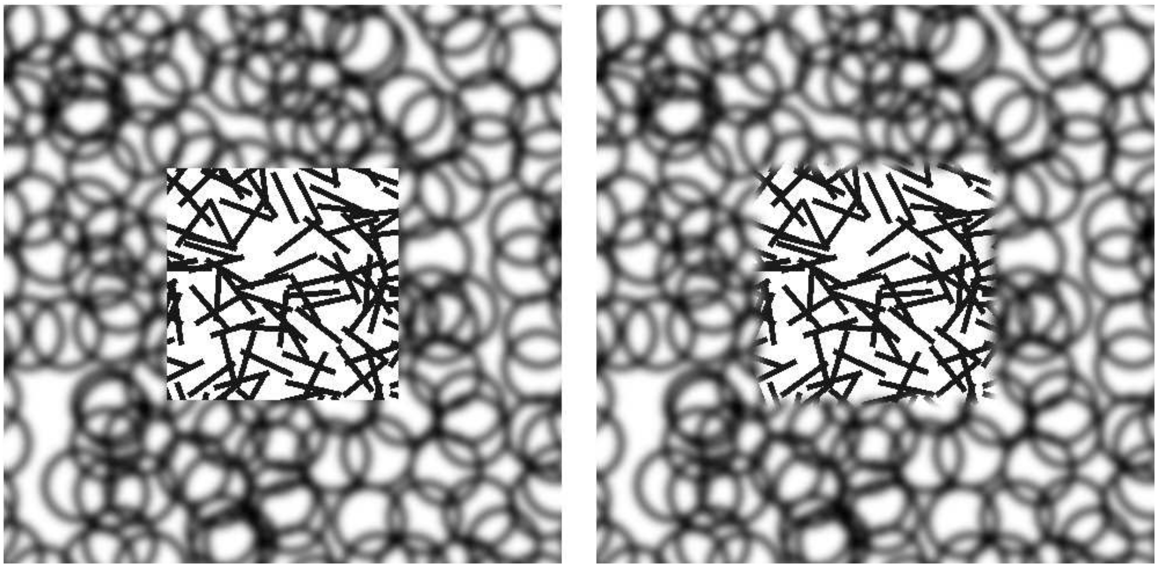

Figure 1: Which central square looks closer? The only difference between the two images is the degree of blur/sharpness of the outer edge of the central square. (A) The image was derived by simulating the lens optics that arise when one focuses on a small line-textured square in front of a circle-textured object. This gives rise to a sharp boundary. (B) The image was derived by simulating the lens optics that arise when one focuses on a line-textured object visible through a small square *hole* in a circle-textured object. This gives rise to an occlusion edge boundary with a blurry appearance.

Consider Fig. 2. The object that is in front determines the characteristics of the boundary.[5] If the object in front is in focus, as illustrated in the left column of Fig. 2, the boundary will also be in focus. If the object in front is out of focus (because the observer is accommodating on the farther region), as illustrated in the right column of Fig. 2, the boundary will be out of focus too (as described in detail below). Thus by noting the relationship between the sharpness of the boundary and the sharpness of the regions to either side of it (assuming they are not uniform), the observer could, in theory, determine which region the boundary "belongs to"[5] and hence which region lies in front. The goal of the study reported here was to determine whether human observers use this information in making depth judgments (e.g., Fig. 1).

We ran two experiments to assess human observers' perception of images with the occlusion edge blur cue. In Experiment 1, we presented pairs of images that differ only in the presence or absence of occlusion edge blur. In Experiment 2, we studied the cue in single images.

Experiment 1: Concentric Squares

Methods

Stimuli

The stimuli in Experiment 1 were 2D gray-scale representations of regions of texture in the configuration shown in Fig. 1. The outer squares subtended $14.7^\circ \times 14.9^\circ$; the inner squares subtended $6.2^\circ \times 6.1^\circ$; the horizontal width of the gap between the outer squares was 0.9°. Viewing was monocular from a distance of 50 cm. The images were displayed on a high-resolution computer monitor.

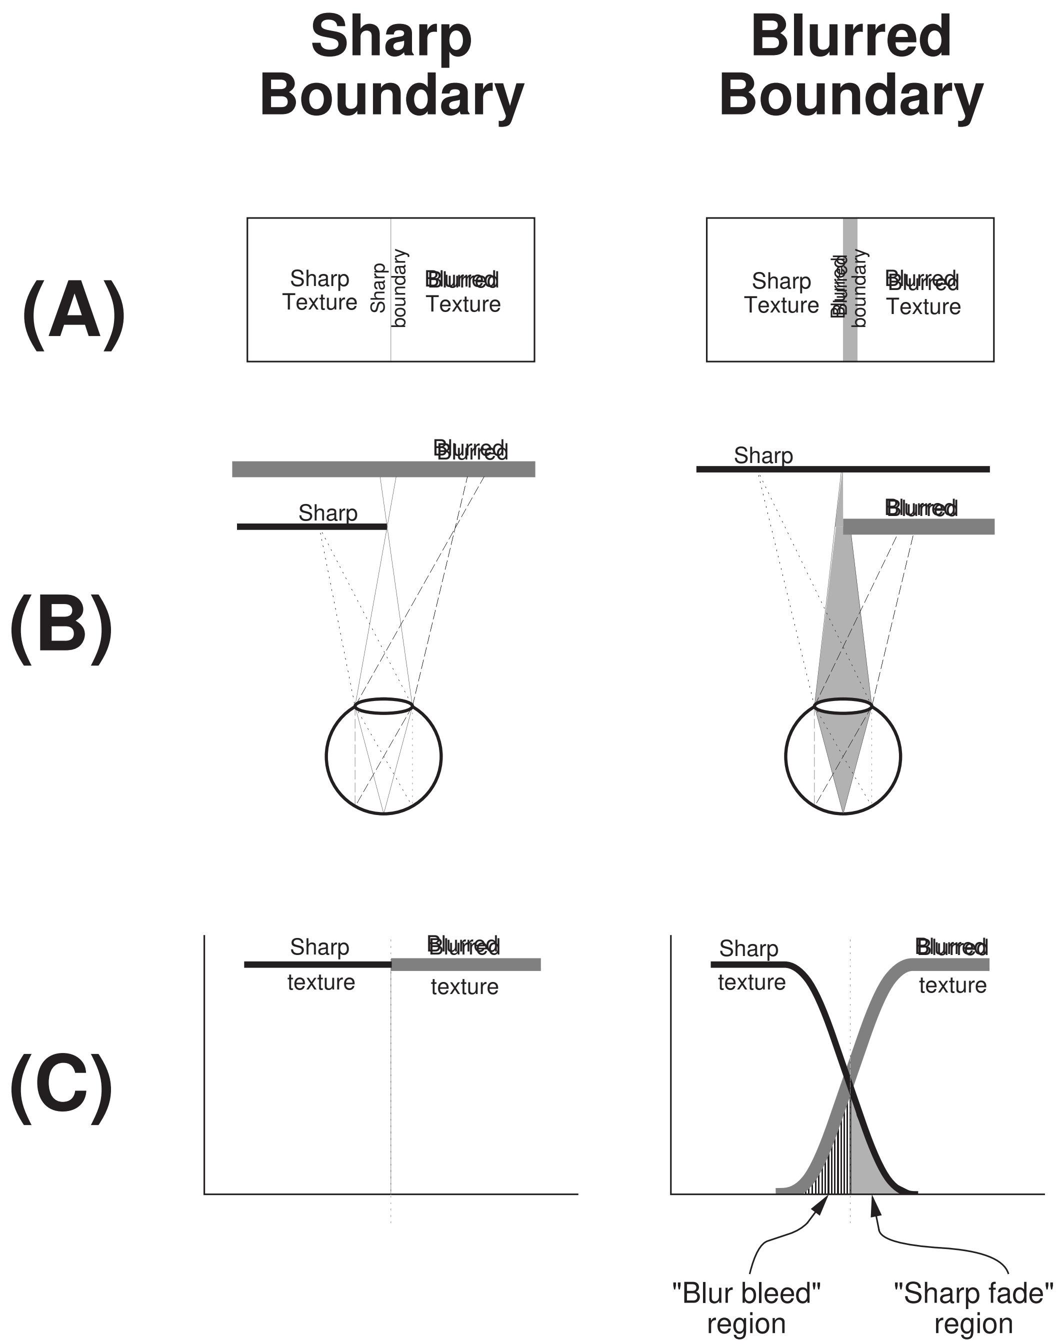

Figure 2: Two kinds of boundaries at a depth discontinuity. If one object at a depth discontinuity is in sharp focus, then the other object is blurred. (A) A sharp optical boundary (left) at the depth discontinuity, and an optical boundary with occlusion edge blur (right) at the depth discontinuity. (B) The sharp boundary (left) is produced when the sharply focused object is *nearer* than the blurred object. The boundary with occlusion edge blur (right) is produced when the sharply focused object is *farther* than the blurred object. The occlusion edge blur arises because image points near the boundary receive a mixture of light rays from both the sharp and the blurred objects. The cone of light rays that converge on a sample image point near the boundary is shown by the shading. (C) The amount of contribution from the sharp and blurred textures is plotted as a function of position across the boundary.

Two stimuli were displayed on a given trial, one in each left/right hemifield. They differed from each other only in the sharpness of the edge dividing the inner and outer squares. One stimulus had a sharp dividing edge and one a blurry dividing edge. The blurriness of the textures themselves was manipulated independently of the edge blur: the outer texture could be blurred (e.g., Fig. 1) or the inner texture could be (e.g., Fig. 3). Four texture patterns were used in all combinations, yielding 32 texture pairs: 4 inner textures $\times$ 4 outer textures $\times$ 2 blur assignments (inner blurred or outer blurred). The two stimuli in an image were presented in random left/right order on the display. The temporal order of presentation of these images was block randomized with a block consisting of one presentation of each of the 32 images. Data were collected from 6 blocks of trials for each observer.

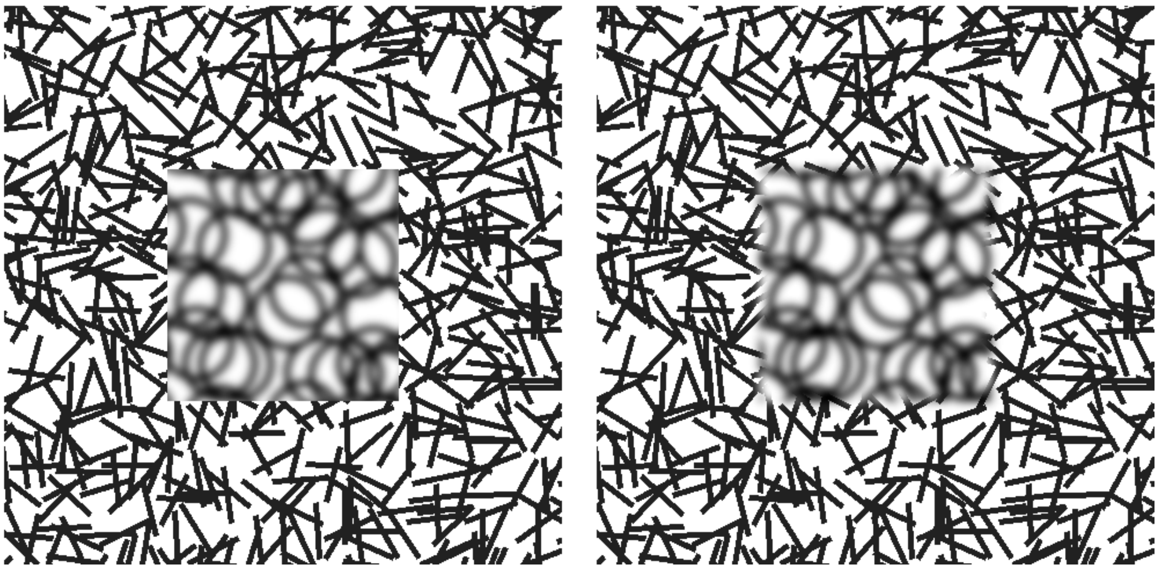

Figure 3: Inner square is blurred.

The blur kernel used in the stimuli for Experiment 1 was a Gaussian with standard deviation $\sigma = 4$. Each pixel measured approximately 106.5 seconds of arc in height and width.

Simulating Occlusion Edge Blur

When two adjacent regions in a 3D scene are located at different depths from the observer, and an observer is accommodating on the near one, as sketched in the left half of Fig. 2, the near texture is sharp, the far one blurred, and the boundary is sharp. This effect is readily simulated in a 2D image by drawing pixels from one (sharp) texture on one side of the boundary and drawing them from the other (blurred) texture on the other side of the boundary. When an observer is accommodating on the far surface, however, a highly specific form of blurring occurs at the boundary between the regions. The boundary characteristics are determined by a combination of the optical blur of the near surface and the vignetting of the far surface near the boundary – the near surface acting as an aperture for the far surface. Optical blur attenuates the high spatial frequency components of the image. Vignetting causes a reduction in luminance (without blurring) of the image of the far surface near the edge. The contributions of both components were simulated in our 2D stimuli. Details of the simulation are given in Appendix A.

Observers

The observers were 11 undergraduate students who were paid to participate in the task. All were naive as to the purpose of the experiment.

Procedure

The observer's task was to report, via a key press, which of the two inner squares – left or right in the stimulus pair – appeared closer in depth. Viewing duration was unlimited and was terminated by the key press. After each trial, the screen was cleared for 5–10 sec as the next image was being prepared for display.

Results

The data were analyzed to determine the percentage of trials in which the observer reported that the inner square was nearer in the sharp-edged stimulus than in the blurry-edged stimulus. Results did not differ systematically across texture pairs, and the data are collapsed across this variable. As anticipated on theoretical grounds, the results depended significantly on whether the inner (center) or outer (surround) texture region was blurred.

Fig. 4 summarizes the results of Experiment 1 for 9 of the 11 observers. The results for the other two observers are considered separately below.

There are clearly two patterns of behavior. Five observers fell in one group, labeled Group A in Fig. 4, and four observers fell in the other group, labeled Group B.

When the texture of the inner square was sharp (as in Fig. 1), all observers consistently saw the sharp-edged inner square as being closer than the blurry-edged inner square that was presented in the other hemifield. The sharp edge on the inner square was consistent with the sharp texture of the inner square, and it was consistent with accommodation on the inner square.

When the inner square texture was blurred as shown in Fig. 3, however, the results for the two groups diverged. Members of Group A saw the sharp-edged inner square as being nearer, whereas members of Group B saw the blurry-edged inner square as being nearer.

Our other two observers for did not behave so consistently. Both judged the sharp-edged inner square to be closer 100% of the time when the inner square texture was sharp, consistent with the other observers, but exhibited inconsistent behavior when the inner square was blurred. One observer judged the sharp-edged blurry inner square to be closer 61% of the time. This ambivalence was fairly consistent across blocks: 50%, 56%, 87%, 31%, 50%, 94% were the block values. The other observer judged the sharp-edged blurry inner square to be closer 72% of the time. This number is misleading, however, because this observer's block data were bimodal: 100%, 100%, 100%, 100%, 12%, and 19%. These results suggest that this observer adopted both of the strategies discussed above, in effect switching from Group A to Group B after the fourth block.

Discussion

The main finding of Experiment 1 is that 9 of the 11 observers consistently used the blur of the texture boundary as a cue to relative depth under the experimental conditions, whether the texture of the inner square was blurred or sharp. The other two observers used this edge-sharpness cue consistently when the inner region was sharp, but not when it was blurred.

Group A's results, that the sharp-edged (blurry) inner square appeared nearer, were not consistent with our original expectation that the dividing edge would be "attached"[5] to a region on the basis of commonality of blur. The texture blur had no discernible effect for these

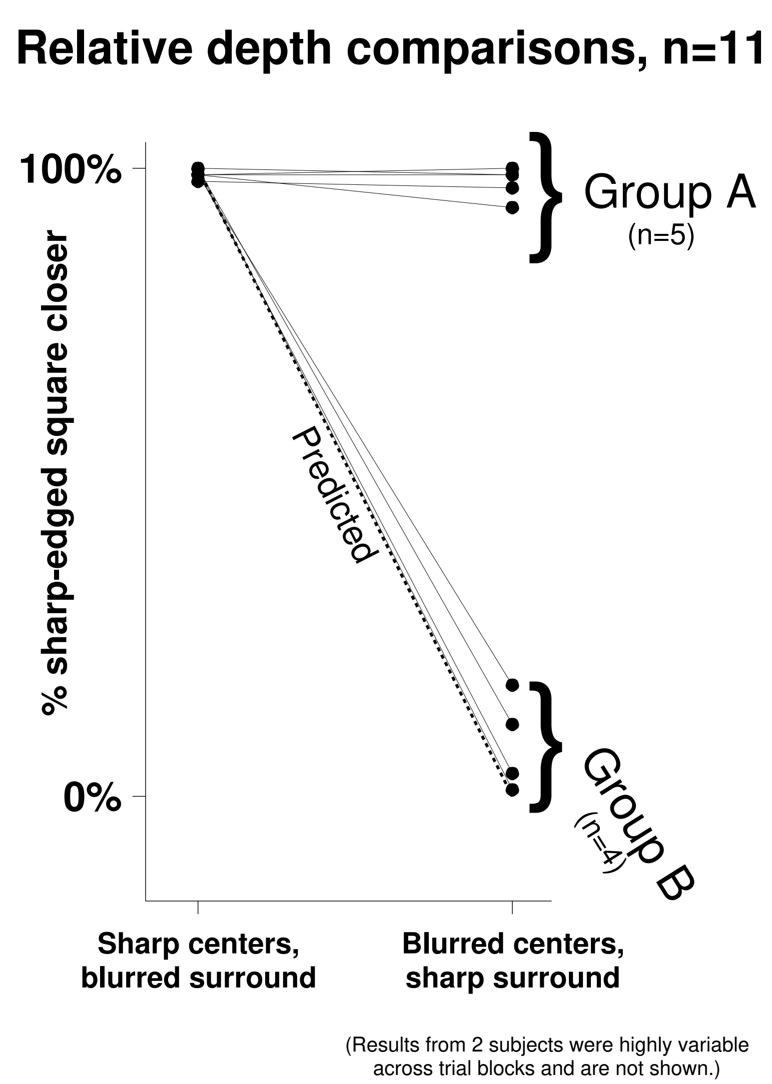

Figure 4: Experiment 1 results. The graph plots the percentage of trials in which each subject judged that the image containing a sharp boundary between the inner and outer squares appeared closer than the image containing a blurred boundary. Two conditions are shown: images where the inner square's texture is in sharp focus and the outer square's texture is blurred (left side of graph); and images where the inner square's texture is blurred and the outer square's texture is sharp (right side of graph).

observers. Instead their results are consistent with the following speculative interpretation, which is based strongly on the geometry of the stimulus. In the absence of other relative depth cues, the smaller inner square would tend to be seen in front of the outer region; this allows both regions to be seen as convex (square) figures. The inner square is also most likely to be the object of fixation, and a near object to which one is accommodated will have a sharp edge. Consequently, the stimulus containing the sharp-edged inner square (Figs. 1A and 3A) is consistent with this interpretation of the geometry, whereas the stimulus containing the blurry-edged inner square (Figs. 1B and 3B) is not. Under this interpretation, the blurred inner texture would be seen as being intrinsically blurred, not blurred because of accommodation on a different depth plane. This interpretation was suggested by the informal reports of some members in this group that both inner squares appeared to be in front of their outer squares, but the sharp-edged one appeared to be more in front.

Group B exhibited very different behavior when the inner texture was blurred, behavior that matched our original expectation. When the inner texture was blurred, members of this group saw the blurry-edged inner square as being more in front than the sharp-edged square. Attachment of the edge to the region with the same blur would cause the (blurry) inner square to be seen in front when the edge was blurry (Fig. 4, right), whereas it would cause the (sharp) outer surface to be seen in front when the edge was sharp (Fig. 4, left). Thus, edge attachment on the basis of commonality of blur could explain the data of Group B.

Although the responses of the members of the two groups differed dramatically from one another, the responses of individual observers within these groups were highly consistent. The data indicate that the stimuli in the two half-images were seen as being unambiguously different in the relative depths of their two texture regions. For these observers, the blur of the inner/outer edge provided a sufficient cue to relative depth in these stimuli. For roughly half of the observers, this cue was linked to the blurriness of the texture regions; for the other half it was independent of texture blur and dependent solely on the spatial arrangement.

Any observer can be biased toward using the occlusion edge blur cue in conjunction either with the inner/outer geometrical relationship (as in Group A) or with the sharp/blurred texture information (as in Group B), by receiving simple instructions on how to form the judgments, or by receiving feedback. Thus, Experiment 1 measured not the usability of the occlusion edge blur cue, but rather the predisposition of naive observers to use the occlusion edge blur cue to make relative depth judgments consistent with the optics of blur in 3D scenes. Four of the 11 observers (those in Group B) were predisposed to make such judgments. Five of the 11 (those in Group A) were predisposed to make judgments using the occlusion edge blur information with the geometric occlusion cue.

Experiment 2: Bipartite Field

Rationale

The concentric squares experiment (Experiment 1) provided compelling evidence of the effect of occlusion edge blur as a cue to relative depth. To test the generality of the cue, we adopted another stimulus and procedure. We chose a simple bipartite field viewed through a highly blurred aperture; the geometry of this stimulus was unbiased.

We also used a different experimental procedure. In Experiment 1, the observer reported which inner square appeared nearer. This left open the possibility that both inner squares could appear to lie in front of (or both behind) the outer squares. In Experiment 2, we used a single bipartite field and asked the observers to make a single judgment of relative depth.

Methods

Stimuli

The stimuli in Experiment 2 were 2D gray-scale representations of regions of texture in the configuration shown in Fig. 5. The $12° \times 12°$ image was divided vertically into two equal-sized texture regions, of which one was sharp and the other blurred. Viewing was monocular from a distance of 50 cm. The images were displayed on a high-resolution computer monitor.

The images were viewed through a small circular aperture in a black opaque occluder positioned near the observer's eye (at a distance of approximately 10–15 cm from the observer). By virtue of the proximity of the occluder, the circular outer edge of the image was highly blurred. Thus, the circular outer edge of the image was seen as belonging to neither texture region.

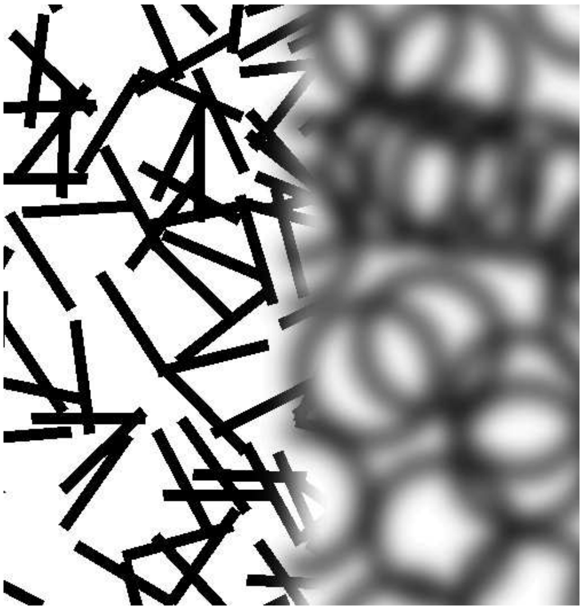

Figure 5: Blurred boundary between two texture regions.

One stimulus was displayed on each trial. In half the trials, the vertical boundary between the two texture regions was sharp; in the other half, it was blurred, using the occlusion edge blur procedure described above. Four texture patterns were used in Experiment 2 – the same patterns as in Experiment 1. Different patterns were chosen randomly for the left and right halves of each display. Thus, there were 48 stimulus images: 4 left textures $\times$ 3 right textures $\times$ 2 blur assignments (left blurred or right blurred) $\times$ 2 edge blur assignments (sharp edge or blurred edge). Each block contained all eight images for two texture pairs; thus a complete test of all images took six blocks. The stimuli were presented in random temporal order within each block. Data were collected from three complete sets of blocks of trials for each observer (144 images per observer), except as indicated in the caption for Fig. 6.

The blur kernel used in the stimuli for Experiment 2 was a Gaussian with standard deviation $\sigma = 12$. Each pixel measured approximately 72 seconds of arc in height and width.

Observers

The observers were five undergraduate students who were paid to participate in the task. All were naive as to the purpose of the experiment.

Procedure

A chin rest was used to stabilize the observer's head. The observer was instructed to adjust the aperture to the largest diameter for which the edges of the image were not visible.

The observer's task was to foveate the center of the circular display and to report, via a key press, which of the two halves – left or right in the stimulus image – appeared closer in depth. Viewing duration was 500 ms. Between trials, the screen was a uniform 72 cd/m^2.

Results

The data were analyzed to determine the percentage of trials in which the observer judged that the blurred side appears closer than the sharp side. Results did not differ systematically across texture pairs, and the data are collapsed across this variable. Fig. 6 summarizes the results for Experiment 2. For all five observers, the blurred side appeared closer more often when the edge was blurred than when it was sharp.

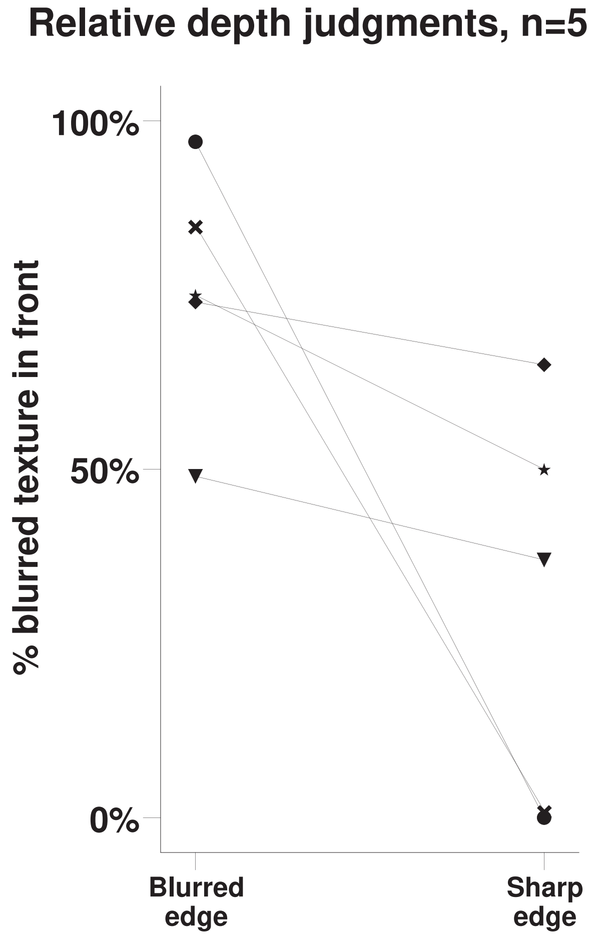

Figure 6: Experiment 2 results. The graph plots the percentage of trials in which each subject judged that the blurred texture region was in front of the sharp texture region. Results are shown for five subjects, on two conditions: images where the boundary between the two texture regions is blurred and images where the boundary is sharp. One observer (★) completed 120 trials; a second observer (▼) completed 136 trials. The other three observers completed 144 trials, as described in the Methods section.

There was also an effect of texture blur that varied across observers: for two of the observers (★ and ◆ in Fig. 6), the blurred half-image tended to be judged as appearing closer more often than did the sharp half-image; for one observer (▼), the sharp half-image appeared slightly closer; and for the other two observers (● and ✖), there was no strong bias.

Although the number of observers was smaller in Experiment 2, it again appeared that the observers fell into two groups, based on the magnitude of the occlusion edge blur effect. Three observers had a relatively weak effect of occlusion edge blur (shallow slopes in Fig. 6); their judgments were strongly biased by the blurriness of the texture itself. The other two observers had a strong effect of occlusion edge blur (steep slopes in Fig. 6) and no apparent bias.

No significant effect of textures or of polarity (left/right) was found.

Discussion

The results of Experiment 2 clearly support our hypothesis that occlusion edge blur acts as a relative depth cue. All five observers judged that the blurred half-image appeared closer more often when the occlusion edge was blurred than when it was sharp.

This experiment provided a direct test of the effect of occlusion edge blur on relative depth judgments. The only other available cue that the observers could use to judge relative depth was the blurriness of the textures themselves. Three of the five observers used this cue; two did not. All of the observers used the edge blur in addition. As in Experiment 1, different observers assigned different weights to the various available depth cues.

General Discussion

The primary finding in the two experiments is that occlusion edge blur is used by human observers as a relative depth cue. Nine of the eleven observers in Experiment 1 and all five observers in Experiment 2 made judgments that depended on occlusion edge blur. However, the results also suggest that different observers assigned different weights to the various available depth cues.[6]

The occlusion edge blur cue requires that the regions being compared differ from each other in blurriness. In addition, the regions must be able to be seen in an occlusion relation (e.g., the left region occluding the right region). Both of these stimulus features can themselves be cues to relative depth. In Experiment 1, we eliminated the effects of the texture blur cue by presenting a pair of images that differed only in the edge blur cue. However, observers had to *compare* the relative depths seen in the two images. In Experiment 2, the observers reported their direct judgment of the relative depths of the two regions in a single image, but the difference in texture blur remained as a potential depth cue (which, in fact, was used by some observers). The two experiments thus complement each other, together showing that occlusion edge blur affects the perceived relative depths of adjoining texture regions.

The Weisstein Effect

In related work, Weisstein and colleagues[7–10] have shown that texturing regions with spatial frequency gratings can affect figure/ground relations. When the faces in a Rubin faces/vase figure[11] are textured with a low spatial frequency grating and the vase is textured with a high spatial frequency grating, the vase appears as figure more often. Conversely, when the faces are textured with the higher spatial frequency grating, they appear as figure more often. This effect is of particular interest because it runs counter to the well-known size cue to visual depth: smaller objects appear to be farther away, all other things being equal. In the Weisstein figures, the smaller texture bars make the region look closer.

Weisstein's effect may be related to the occlusion edge blur cue. The border in the faces/vase figure is sharp. If the higher spatial frequency is analogous to our sharp texture and the lower spatial frequency analogous to our blurred texture, then the occlusion edge blur cue would suggest that the higher spatial frequency region appears to be in front because the border is seen to be attached to it. If this explanation is correct, blurring the border in the faces/vase figure (in the manner described in Appendix A) should cause the Weisstein effect to reverse, or at least to become observer-dependent.

Occlusion Edge Blur Differs from Other Depth Cues

Occlusion edge blur differs from the oculomotor cue of accommodation,[1,2] in that no active change in eye focus is required. Relative depth information can be inferred from occlusion edge blur in a static image. Occlusion edge blur also differs from pictorial blur (or haze blur),[1,2] which exists in images of very distant objects (e.g., mountains on the horizon), affecting all distant visual features, regardless of whether they are depth discontinuities.

Occlusion Edge Blur in Image Display

The occlusion edge blur depth cue is usually absent in graphic image displays; all image points are displayed in sharp focus. The absence of this depth cue may cause displayed images to appear flatter (less 3-dimensional) than they otherwise would,[6] but implementing blur in graphic image displays is not simple.

To use the cue in a static image, one has to select a depth in the 3D scene to be the focal plane. If the scene being represented has a natural focus of attention, this may be done in a reasonable way. Displaying an image with only the chosen plane in focus might serve to direct the observers' attention to information at that depth. In a dynamic system, real-time measurements of the observer's eye position and accommodation could be used to enable presentation of an image in which the blurriness of the various regions corresponds more accurately to the 3D percept.

The following conditions were found to be sufficient to evoke a response to the edge blur cue. They may not all be necessary.

- *Textured Objects* – The regions on both sides of the boundary must have some visible texture to enable the observers to determine whether the regions are blurred or sharp.
- *Intrinsic Sharpness* – Our analysis assumes that the textures on the objects depicted in the image are intrinsically sharp. Few real object textures are physically blurred. An exception might be certain kinds of stone, where veins of darker and lighter material have melted together to form gradual ("blurred") transitions. For such images, erroneous interpretations might arise. Group A in our study may have seen the blurred texture as being intrinsically blurred.
- *Sufficient Blur* – The depth difference at the discontinuity must be large enough to produce detectable blur.

Occlusion Edge Blur as a Means of Resolving Depth-from-Focus Ambiguity

It is well known that the degree of blur of a visual object serves as a coarse depth cue: the blurrier the object, the farther or closer it is in depth from the plane of focus. If the focal distance of the lens is known, then it can be used to scale the blur measurement, producing an absolute depth estimate. Several researchers have proposed methods by which measurements of the degree of blur in parts of an image can be used to compute depth information.[12–19]

One cannot, however, determine absolute depth from single-image focus information alone. If the image of an object is defocused, it may be so because the object is closer than the plane of focus or because the object is farther than the plane of focus. Therefore, there is a near/far ambiguity in depth-from-focus computations.[4,17] To resolve this ambiguity, computional algorithms for computing depth-from-focus have introduced ad hoc assumptions. For example, the assumption has been made that images are focused on the nearest object[17] and consequently that the blurred objects are behind the focal plane.

Such assumptions are unrealistic in human vision, however, which does not follow such simple rules.

The occlusion edge blur cue may provide a more realistic way to resolve the ambiguity in cases where there are adjacent blurred and sharp regions in an image. The sign of depth of the blurred region relative to the sharp region can be determined by examining the boundary between the two regions. If the boundary has the characteristic blur/vignetting structure of occlusion edge blur, then the blurred region can be judged to be in front of the sharp region. If the boundary is sharp, then the sharp region can be judged to be in front of the blurred region.

Conclusions

Occlusion edge blur has been shown to be an effective cue for relative depth. Its weighting, in comparison to other depth cues, varies across observers. In sufficiently structured images, the cue resolves the near/far ambiguity inherent in depth-from-focus computations.

Acknowledgments

Supported in part by NIH (P01 CA47982-04), AFOSR (F49620-92-J-0114), Oak Ridge Associated Universities (ORAU Junior Faculty Enhancement Award), and UNC-CH (Junior Faculty Development Award). We thank Gal Zauberman and Xiaofei Wang for valuable assistance with the experiments.

Some of the results in this paper were previously reported at the 1992 Annual Meeting of the Association for Research in Vision and Ophthalmology (ARVO).[20]

Appendix: Calculations for Occlusion Edge Blur Images

In the stimuli with a blurry dividing edge, the intensity I of each image point was formed according to the equation

$$I \;=\; G * B \;+\; (G * B') \cdot S.$$

G is the Gaussian blur kernel described by $G(x,y) \;=\; \frac{1}{\sigma\sqrt{2\pi}} e^{-(x^2+y^2)/2\sigma^2}$, where x and y are measured in units of pixels on the display, and σ is the standard deviation of the Gaussian. B is the intensity image of the (occluding) blurred object, B' is a step-function image that is 0 everywhere across the blurred object and 1 elsewhere, S is the intensity image of the (occluded) sharp object, and $*$ denotes the convolution operation. A Gaussian was chosen to model the blur kernel because that function is a reasonable first approximation to the blur in human vision.[17]

The $G * B$ term describes how the blurred object contributes to the image (Fig. 2A, right column). This term can be understood by imagining a scene with an occluding surface B in front of a totally black (zero intensity) occluded background surface. Since the occluder is blurred in this case, the image would be described by $G * B$ (Fig. 7, left).

The contribution from the sharp occluded surface is described by the $(G * B') \cdot S$ term. This characterizes the vignetting of the image of the occluded surface (Fig. 2B). If both the occluding and occluded surfaces were sharp (i.e., at the same depth), then G would be a delta function, and this term would reduce to $B' \cdot S$, which (by the definition of B') is just the visible (unoccluded) part of the occluded surface, and is therefore correct in this case. But in the more general case where the occluder is blurred, there is a gradual transition band from the regions where the rear surface is completely visible to the regions where it is completely invisible. At any point in the transition band, the amount of light is

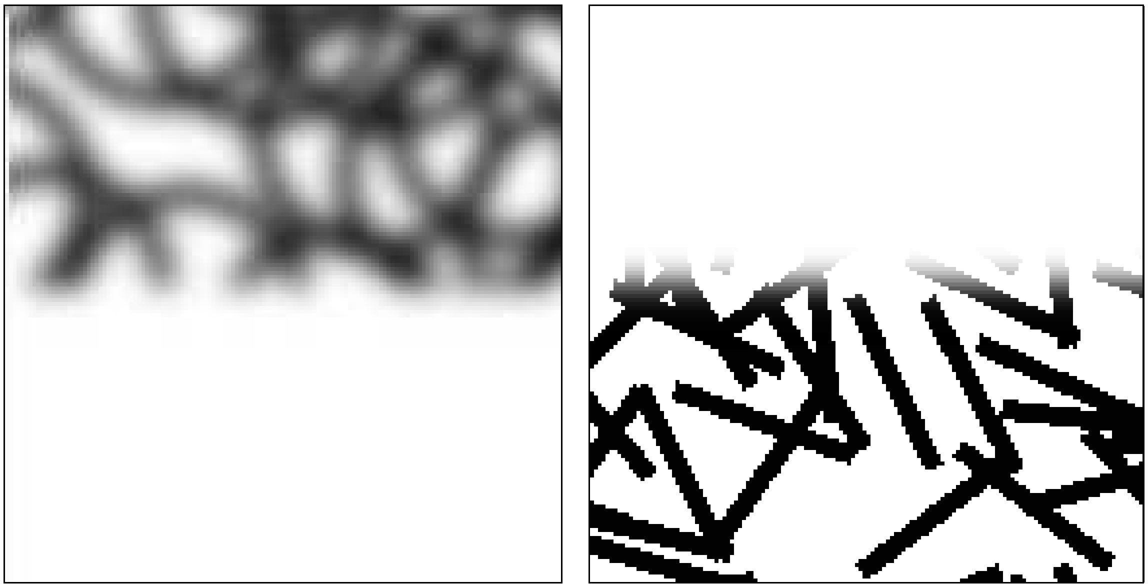

Figure 7: Detail of blur and vignetting at boundary. (Left) Contribution to the image from a near blurred object. (Right) Contribution to the image from a far sharp object.

equal to the image intensity at that point, attenuated in proportion to the fraction of rays occluded by the near surface. Since the rays that can strike the image point radiate from the scene point in proportion to the blur kernel, the fraction of rays *not* occluded is equal to the blur kernel times a step function that is 0 everywhere across the blurred object and 1 elsewhere. This vignetting function $G * B'$ describes the fraction of the original scene point intensity S reaching the image point. The product $(G * B') \cdot S$ thus describes the contribution from the sharp surface to the image, which fades in intensity, as shown in Fig. 7 (right). (If the occluder were semi-transparent, then B' would have a value between 0 and 1 everywhere across the blurred object and 1 elsewhere.)

The total amount of light reaching an image point is then just the sum of the amounts from the near surface and from the far surface; thus, $I = G * B + (G * B') \cdot S$.

References

1. V. Bruce and P.R. Green, *Visual Perception: Physiology, Psychology, and Ecology*, 2nd edition (Lawrence Erlbaum Associates, Hillsdale, NJ, 1990).
2. L.R. Wanger, J.A. Ferwerda, and D.P. Greenberg, "Perceiving spatial relationships in computer-generated images," IEEE Computer Graphics & Applications, **11**, 44–58 (May 1992).
3. A.P. Pentland, "The focal gradient: Optics ecologically salient," Investigative Ophthalmology and Visual Science **26**, Supplement, 243 (1985).
4. P. Grossman, "Depth from focus," Pattern Recognition Letters **5**, 63–69 (1987).
5. K. Nakayama, S. Shimojo, and G.H. Silverman, "Stereoscopic depth: Its relation to image segmentation, grouping, and the recognition of occluded objects," Perception **18**, 55–68 (1989).

6. M.S. Landy, L.T. Maloney, E.B. Johnston, and M. Young, "Measurement and modeling of depth cue combination: In defense of weak fusion," Vision Research **35**, 389–412 (1995).
7. V. Klymenko and N. Weisstein, "Spatial frequency differences can determine figure-ground organization," Journal of Experimental Psychology: Human Perception and Performance **12**, 324–330 (1986).
8. N. Weisstein and E. Wong, "Figure-ground organization and the spatial and temporal responses of the visual system," in *Pattern Recognition by Humans and Machines: Visual Perception* **2**, E.C. Schwab and H.C. Nusbaum, eds. (Academic Press, Orlando, FL, 1986), pp. 31–64.
9. E. Wong and N. Weisstein, "Sharp targets are detected better against a figure and blurred targets are detected better against a ground," Journal of Experimental Psychology: Human Perception and Performance **9**, 194–202 (1983).
10. E. Wong and N. Weisstein, "The effects of perceived figure and ground on the orientation discrimination of sharp and blurred targets," Perception **13** (1984).
11. E. Rubin, "Figure and ground," reprinted in *Readings in perception*, D.C. Beardslee and M. Wertheimer, eds. (Van Nostrand, Princeton, NJ, 1958), pp. 194–203 (1921).
12. T. Darrell and K. Wohn, "Pyramid based depth from focus," IEEE International Conference on Computer Vision and Pattern Recognition (Ann Arbor, MI, 1988), pp. 504–509.
13. J. Ens and P. Lawrence, "An investigation of methods for determining depth from focus," IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-15**, 97–108 (1993).
14. B.K.P. Horn, "Focusing," Artificial Intelligence Memo 160, Massachusetts Institute of Technology.
15. R.A. Jarvis, "A perspective on range finding techniques for computer vision," IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-5**, 122–139 (1983).
16. S.-H. Lai, C.W. Fu, and S. Chang, "A generalized depth estimation algorithm with a single image," IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-14**, 405–411 (1992).
17. A.P. Pentland, "A new sense for depth of field," IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-9**, 523–531 (1987).
18. M. Subbarao, "Parallel depth recovery by changing camera parameters," Proceedings of the IEEE 2nd International Conference on Computer Vision (IEEE Computer Society Press, Washington, DC, 1988), pp. 149–155.
19. M. Subbarao and N. Gurumoorthy, "Depth recovery from blurred edges," IEEE International Conference on Computer Vision and Pattern Recognition (Ann Arbor, MI, 1988), pp. 498–503.
20. J.A. Marshall, K.E. Martin, D. Ariely, C.A. Burbeck, and J.P. Rolland, "Blur attachment as a visual depth cue." Investigative Ophthalmology & Visual Science, **33/4**, 1369 (1992).