

Vertical position as a cue to pictorial depth: Height in the picture plane versus distance to the horizon

JONATHAN S. GARDNER, JOSEPH L. AUSTERWEIL, AND STEPHEN E. PALMER
University of California, Berkeley, California

Two often cited but frequently confused pictorial cues to perceived depth are height in the picture plane (HPP) and distance to the horizon (DH). We report two psychophysical experiments that disentangled their influence on perception of relative depth in pictures of the interior of a schematic room. Experiment 1 showed that when HPP and DH varied independently with both a ceiling and a floor plane visible in the picture, DH alone determined judgments of relative depth; HPP was irrelevant. Experiment 2 studied relative depth perception in single-plane displays (floor only or ceiling only) in which the horizon either was not visible or was always at the midpoint of the target object. When the target object was viewed against either a floor or a ceiling plane, some observers used DH, but others (erroneously) used HPP. In general, when DH is defined and unambiguous, observers use it to determine the relative distance to objects, but when DH is undefined and/or ambiguous, at least some observers use HPP.

Depth cues are sources of information about the distance to different objects and surfaces in the visible environment, relative to the observer's viewpoint. Pictorial depth cues are the subset of depth cues that are available in a single projected image at a particular moment in time (i.e., static, monocular, image-based cues). Many pictorial depth cues are well-known and standardly described in textbook discussions of depth perception (e.g., Palmer, 1999). Two putative pictorial depth cues that are frequently mentioned, but not always clearly differentiated and/or explained (e.g., Blake & Sekuler, 2006; Wolfe et al., 2006), are the subject of this article: height in the picture plane (HPP) and distance to the horizon (DH). Despite their similarity when the center of projection is above the supporting plane, they are conceptually distinct and can be dissociated empirically when the depicted object lies on a flat surface that can be either below or above the horizon. In the present article, we clarify the nature of these two cues and the conditions under which they are used by observers.

As we will use the term, HPP refers to the distance from the picture's lower border to the bottom of the object. Its influence is usually demonstrated in images in which the horizon is implicitly or explicitly above the objects in question, so that the objects are perceived as lying on a horizontal plane below eye level. (The *horizon* of any plane is the set of all vanishing points for parallel lines that lie on that plane.) Its perceptual relevance is that, under such circumstances, objects that are higher in the picture plane are generally seen as more distant from the viewer

(e.g., Figure 1A). DH refers to the distance from the horizon to the point of contact between the object and its supporting surface—that is, the bottom of an object that is supported by a plane below eye level (Figure 1B) or the top of an object that is supported by a plane above eye level (Figure 1C). The influence of DH is usually demonstrated when the horizon is explicitly present in the display. Its perceptual relevance is that objects that are closer to the horizon are usually seen as more distant from the viewer for objects attached to planes both below the horizon (Figure 1B) and above it (Figure 1C).

HPP and DH are completely correlated when the objects in question lie on a horizontal plane below eye level, as when one views objects resting on the ground, a table, or a counter (see Figures 1A and 1B). However, HPP and DH make opposite predictions when the objects are attached to a horizontal plane above eye level, as when one views things attached to a ceiling or some other surface above eye level (see Figures 1B vs. 1C). The geometrically correct answer is always given by DH; HPP is correct only when the objects lie on a plane below the horizon. Furthermore, DH is much more general than HPP, in the sense that it is applicable not only to horizontal planes, but also to any flat surface to which an object is attached, because every flat plane that is not perpendicular to the line of sight has a corresponding "horizon" defined by its vanishing points (Sedgwick, 1986). Nevertheless, it is possible that observers may use HPP heuristically in some ambiguous and/or unfamiliar situations when DH actually provides more accurate information. In the following

S. E. Palmer, palmer@cogsci.berkeley.edu

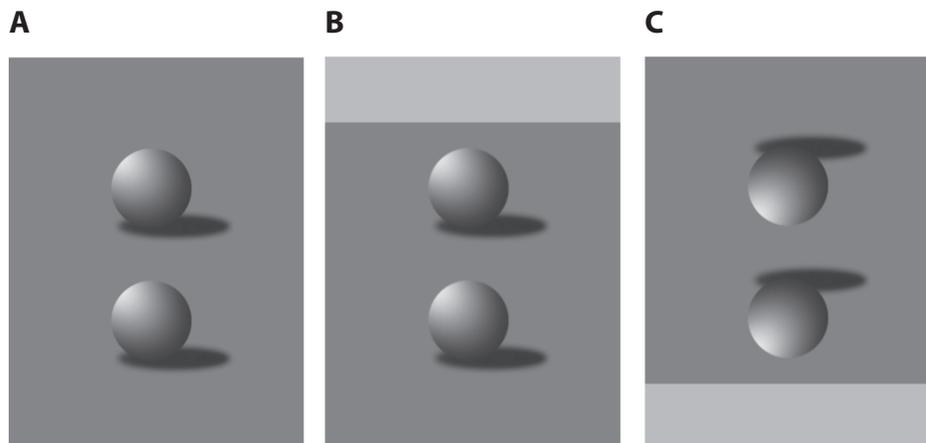


Figure 1. Demonstrations of relative depth perception based on vertical position. The sphere that is higher in the picture plane is perceived as farther away if it is seen as supported by an underhead plane (A and B) but as closer if it is seen as supported by an overhead plane (C).

experiments, we sought to determine the extent to which observers' depth perceptions conform to the predictions of DH versus HPP under different display conditions.

EXPERIMENT 1 Depth Perception in an Enclosed Rectangular Space

To study the effects of HPP and DH on depth perception in displays containing explicit information about the horizon for surfaces both below and above eye level, we independently manipulated the vertical position of a sphere and the vertical position of the horizon within the picture frame over a sample of possible combinations consistent with attachment to the adjacent horizontal plane. Varying the position of both the horizon and the target sphere should also allow us to differentiate between two different formulations of DH, one based on the absolute projected distance to the horizon and the other based on the proportional projected distance to the horizon relative to the relevant edge of the frame. Cast shadows were used to specify whether the sphere was resting on a lower plane (the "floor") or attached to a higher plane (the "ceiling") of a rectangular, three-dimensionally bounded space (the "room") for which the implied horizon of both horizontal surfaces was midway between the edges formed by the back wall with the floor and ceiling.

To generate these scenes, we first used a 3-D modeling program to render a square image of a canonical rectangular room in one-point perspective from a viewpoint midway between the floor and ceiling (see Figure 2A). In order to maximize the appearance of depth, the images included a lightness gradient along the horizontal planes consistent with a light source behind the viewer (Epstein, 1966). We then defined six equally spaced vertical positions within a 4×3 cropping frame at which the target sphere and the horizon could be located. We then translated the cropping frame vertically relative to the room and its horizon to generate four images of the same room in which the hori-

zon was located at Positions H2, H3, H4, and H5 (dashed lines in Figure 2A). A sphere was rendered separately and placed at one of a set of vertically displaced positions physically consistent with attachment to the floor or ceiling plane, using a cast shadow to indicate the appropriate attachment (Mamassian, Knill, & Kersten, 1998). As is shown in Figure 2B, there were four possible object positions (O1–O4) for an object on the ground plane when the horizon was at the highest position (H5), there were three possible object positions (O1–O3) when the horizon was at the next highest position (H4), and so on. At the lowest horizon position (H2), there was only one position for the object (O1), because the object at higher positions could no longer be resting on the ground plane. The images with the sphere on the ceiling plane were simply the ground-plane images reflected vertically about a horizontal axis, thus locating the sphere at Positions 3–6 (see Figure 2C). The sphere and shadow were placed in the image after rendering (rather than being rendered directly in the 3-D modeling program), so that they had the same image-based size (in pixels) at each position, thus eliminating relative object size as a confounding variable.

We created these images by cropping a single larger image, as one might do with a photograph, in order to isolate the effects of HPP and DH as completely as possible from all other image-based features: Nothing other than the vertical position of the objects, surfaces, and horizon in the scene varied. These cropping transformations are not the same as the sorts of transformations that cause differences in the vertical positions of objects and horizons under natural viewing conditions, however. Vertical displacement of objects occurs most often under natural viewing conditions when observers change the direction of their gaze by rotating their eyes and/or head upward or downward (i.e., varying the pitch of the direction of gaze; see, e.g., Ooi, Wu, & He, 2001). However, this rotational transformation, unlike cropping, also changes the slants of horizontal surfaces relative to the direction of gaze. Vertical displacements also occur naturally when

A. Horizon Heights

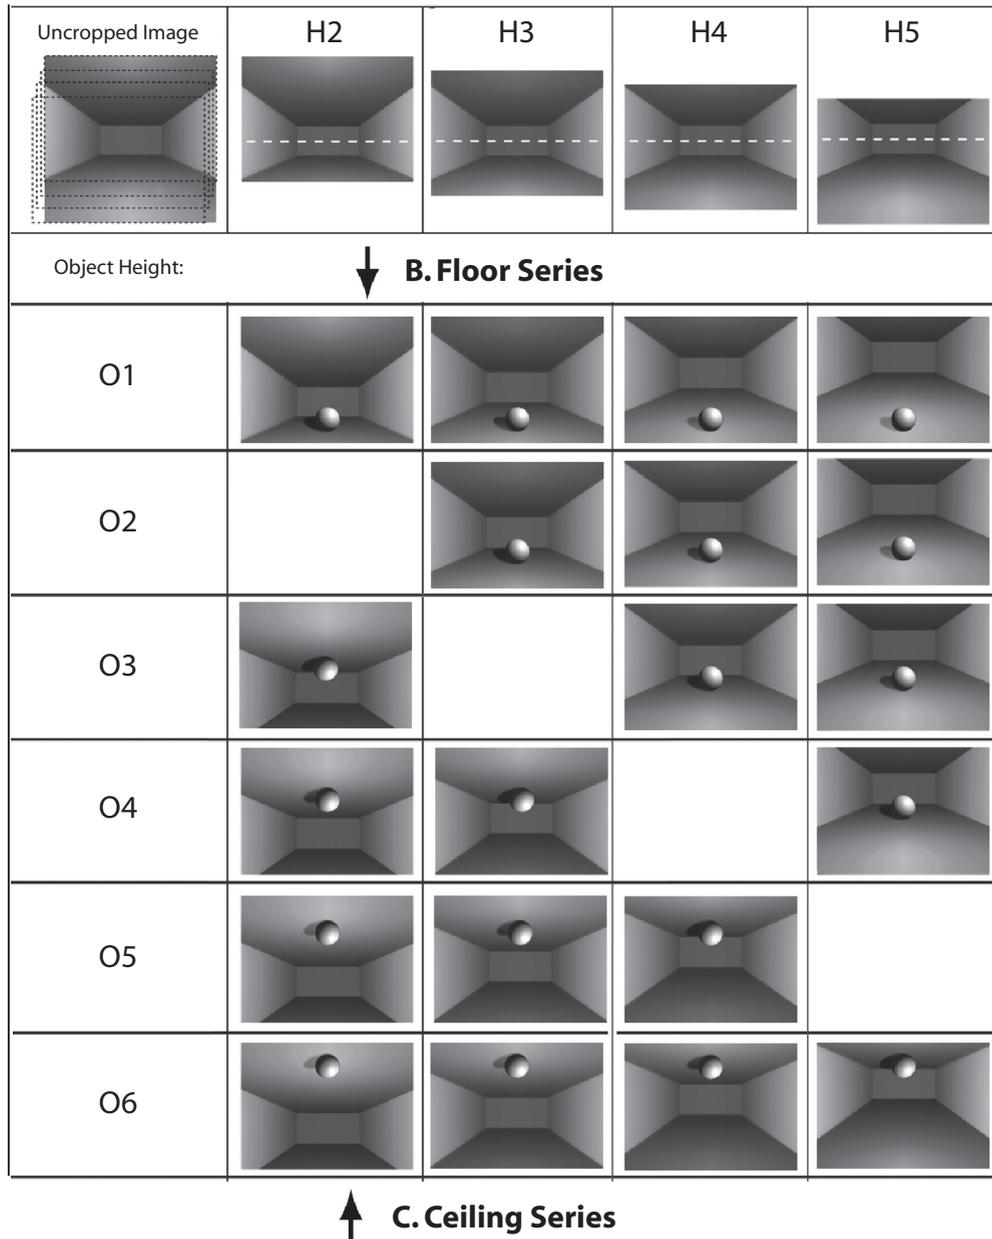


Figure 2. Display construction for Experiment 1. (A) Basic room rendering on which all the images were based and four croppings that resulted in the four horizon height positions (H2, H3, H4, and H5). (B) Addition of the target sphere at each object height position from O1 to O4 to create the set of underhead plane (“floor”) images. The heights of the horizon line and the sphere in the frame are labeled along the rows and columns, respectively. (C) Addition of the target sphere at object heights of O3 to O6 to create the set of overhead plane (“ceiling”) images. Note that for each floor image, the horizon must be higher than the sphere and that the reverse is true in the case of the ceiling images.

observers change the position of their eyes and/or heads by translating them up or down (i.e., varying the height of the center of projection). However, this translational transformation, unlike cropping, also changes the distances of the horizontal surfaces from the center of projection. We do not believe that these differences present a major problem, because the present displays *look* as if the slants

and/or distances of the floor and ceiling planes differ in the different croppings of the single larger image. Strictly speaking, however, the effects we report below apply to the perception of depth in cropped photographs and realistic paintings but may have strong implications for perspective changes produced by vertical eye–head rotations and/or translations.

Method

Participants. All 11 participants were students at the University of California, Berkeley, who received partial course credit in their undergraduate psychology course. Their mean age was 20.2 years. All had normal or corrected-to-normal vision, were naive as to the purpose and nature of the experiment, and gave informed consent in accord with the policies of the University of California, Berkeley, Committee for the Protection of Human Subjects, which approved the experimental protocol.

Design. There were 380 paired comparison displays that consisted of all possible ordered pairs of the 20 images defined above (Figure 2B): 10 in which the sphere was attached to the ground plane (all four positions at the highest horizon position, three at the next position, then two, then one) plus the corresponding 10 in which it was attached to the ceiling plane. The order of these trials was randomized by Presentation software (www.neurobs.com), which controlled the display of images and the collection of responses.

Displays. All the scenes were initially modeled and rendered using Carrara Studio 5 software (www.daz3d.com), and the resulting images and screens for the experiment were constructed using Adobe Photoshop CS3 (www.adobe.com). Each image measured 420×315 pixels, and the positions of the object and the horizon line were 44 pixels apart (e.g., both the sphere and horizon line at O2 were 44 pixels higher than their respective positions at O1). The sphere was 85 pixels in diameter and was colored gray, with shading to specify its curvature. All four surfaces of the room were grayscale as well (Figure 2). The spheres were the same in projected size in each image. On each trial, two vertically centered pictures were presented side by side, with 130 pixels separating them. The monitor was an Iiyama Vision Master 451 color monitor that measured 18 in. diagonally, and the resolution was $1,024 \times 768$ pixels, with a refresh rate of 85 Hz.

Procedure. The participants viewed the computer screen from approximately 60 cm. They were instructed to look at each screen and to press a button (left or right) to indicate the image in which the spherical object appeared closer to themselves. Each pair of images remained on the screen until a response was made.

Results and Discussion

We computed the average probability of choosing each image as depicting the farther sphere in all of its pairwise comparisons to provide an overall measure of the relative perceived distance of the sphere from the viewer in each image. The resulting probabilities, averaged over participants, are plotted in Figure 3 as a function of the distance of the sphere from the horizon (DH) and the distance of the horizon from the relevant edge (i.e., the lower edge of the frame if the sphere was on the floor and the upper edge of the frame if it was on the ceiling). The data from the floor conditions are plotted in Figure 3A, and those from the ceiling conditions in Figure 3B, although each data point includes choices from comparisons containing both floor and ceiling images.

The results show that this probabilistic measure of perceived distance is an inverse linear function of DH: The sphere looks farthest from the observer when it is closest to the horizon and closest to the observer when it is farthest from the horizon. This pattern of results is fully consistent with the use of DH to judge depth, but not with the use of HPP. Indeed, HPP alone has a correlation with the present results of only .07, accounting for just 0.5% of the variance. In contrast, DH alone has a correlation of .83 with the results, accounting for nearly 70% of the variance. Clearly, the present findings imply that DH, rather

than HPP, is the more relevant cue in perceiving the distance to an object, at least under conditions in which both floor and ceiling planes are visible within the picture.

Nevertheless, factors other than absolute DH clearly influenced the results. Consider first the floor conditions (Figure 3A). The functions for the different horizon positions do not lie on top of each other, as they would if observers had perceived the depth of the sphere entirely in terms of its absolute (projected) distance from the horizon. Rather, the slopes of the DH functions fan out systematically according to the position of the horizon within the frame, flattening as the horizon is located farther from the closest visible edge of the relevant surface. This result suggests that observers judged the depth of the sphere in terms of *proportional* DH rather than *absolute* DH. That is, they appear to have perceived the sphere as lying at a distance that is an inverse linear function of the (projected) distance from the attachment point of the sphere to the horizon, relative to the total (projected) distance from the closest visible edge of that surface to its horizon. (Note that the horizon specified here is the true horizon of the relevant surface, rather than its farther visible edge at the back of the room.)

To make this proportional DH pattern clear, Figure 3A plots the implicit position of the true horizon (not the far edge of the floor) by extrapolating the data curves to the open square at the best-fitting y -intercept (zero point) and the implicit position of the lower edge of the picture by extrapolating the data curves for each horizon condition to the appropriate proportional distance from the horizon, relative to the sphere positions at that horizon height. For example, because the point labeled O1 on Curve H2 in Figure 3A corresponds to an image in which the sphere is positioned halfway between the true horizon and the lower edge of the picture, Point O1 on Curve H2 should be located halfway between the y -intercept (representing the distance of the horizon) and the extrapolated solid square (representing the distance to the lower edge of the picture). Point O2 on Curve H4 should likewise be midway between the position of the horizon (open square) at the y -intercept and the position of the lower edge point (solid square), because the sphere in the corresponding image is halfway between the horizon and the lower edge of the picture. The data points for all the conditions fall very close to the straight lines connecting these points, and the horizon is extrapolated to a value very near the theoretical ideal point of 100%, indicating the location at which the sphere would be infinitely far away. This proportional DH model accounts for 99.8% of the variance for the floor data with just two parameters: one that specifies the y -intercept of the horizon and the other that specifies the y -value of the closest edge (Figure 3A).

The data for the ceiling conditions are qualitatively similar but differ in several important respects. First, the y -value of the horizon (at zero DH) is slightly less than those in the floor conditions (about 90%), suggesting that the implicit horizon of the ceiling appeared slightly closer than the implicit horizon of the floor to our observers. Second, the closest edge of the ceiling does not appear to

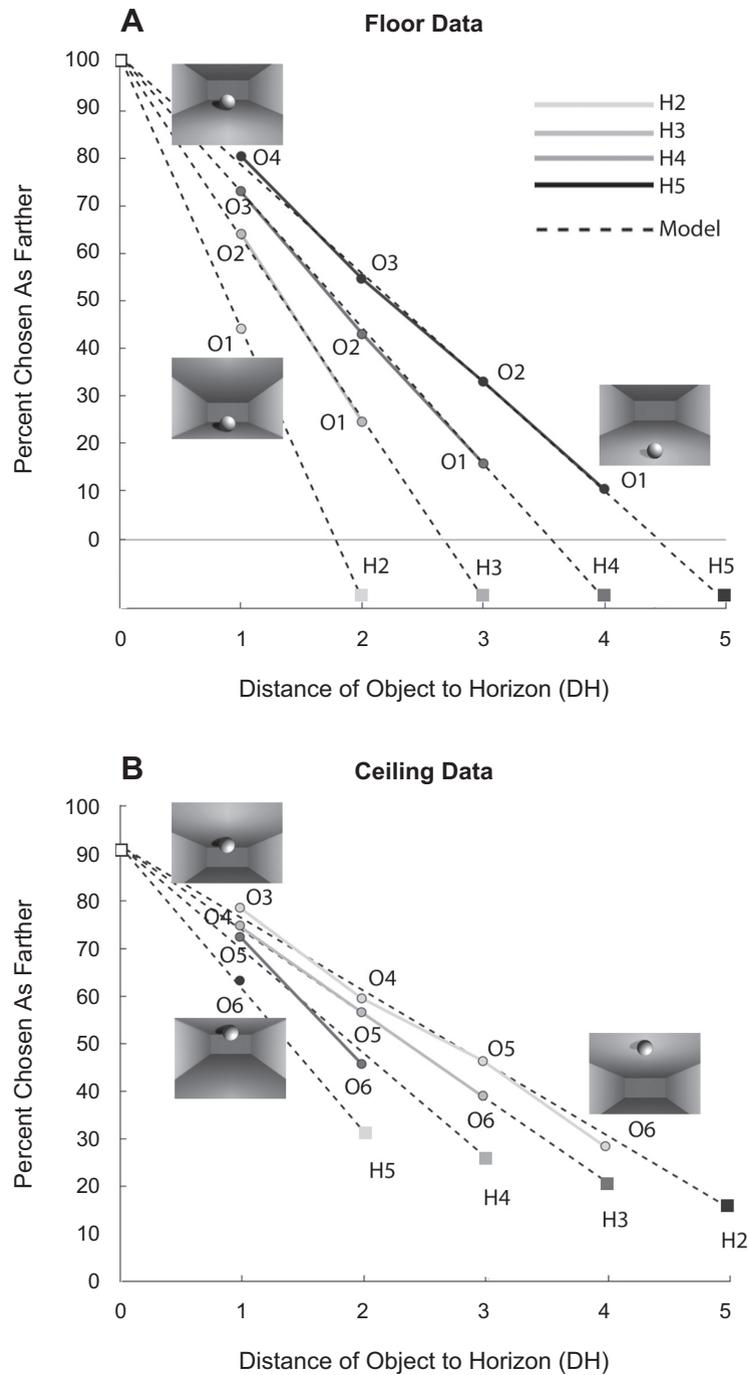


Figure 3. Results of Experiment 1. The circles plot the average percentages of trials on which each image was chosen as farther away, when compared with all other images in pairwise comparisons. (A) Results for the images in which the sphere was attached to the floor. (B) Results for the images in which the sphere was attached to the ceiling. The squares and dashed lines show model fits, as described in the text.

be located at quite the same distance from the observer in the different horizon conditions but, rather, appears to be closer to the observer as the distance of the horizon from the reference (top) edge increases (i.e., closest in the H2 conditions). This result necessarily reduces the effect of

the distance of the horizon from the relevant edge, such that the curves for the different room contexts lie more nearly on top of each other, as they would if an absolute DH model were used. We therefore modeled the ceiling data as a compromise *averaging* model between a purely

proportional DH model and a purely absolute DH model by adding a third parameter that represented the relative weighting of the absolute and proportional DH models. The best fit to the data occurred when the proportional DH predictions were weighted twice as heavily as the absolute DH predictions (66% proportional DH; 34% absolute DH). This averaging model accounts for 99.5% of the variance with three parameters (see Figure 3B). Naturally, the same averaging model fits the floor data for the special case in which the proportional DH component has a weight of one and the absolute DH component has a weight of zero.

The third—and in many ways, most obvious—difference between the floor and ceiling data is that the spheres on the ceiling appeared to be systematically farther away than the spheres on the corresponding floor positions, even though the center of projection of the rendered scene was equidistant from the floor and ceiling planes. When the average choice probabilities are compared for the 20 images (see Figure 3), all 10 of the average choice probabilities for the ceiling points are greater than all 10 of the choice probabilities for the corresponding floor points ($p < .001$ by a sign test). This is also true for all four extrapolated points representing the closest visible points on the ceiling, relative to the closest visible points on the floor. The only case for which the ceiling point was estimated to be closer than the corresponding data floor point is that for the extrapolated horizons. This finding is consistent with a report by Dilda, Creem-Regehr, and Thompson (2005) that people systematically overestimate the distance to ceiling targets.

Are objects on ceiling planes above eye level actually farther from human observers than those on planes below eye level? If so, the bias we have observed would reflect a veridical statistical fact about the observers' environment (a Bayesian *prior*) that influences observers to see objects on overhead planes as farther away than ones at corresponding positions on "underhead" planes. In the absence of relevant ecological data, we cannot know with certainty, but with various assumptions, its plausibility can be evaluated. Standard ceilings are 8 ft high, and the vast majority of adults are between 5 and 6.5 ft tall. If one takes the relevant data to be the actual distances to the floor and ceiling in a standard room, relative to a standing observer's eyes, a *farther ceiling* bias would not be veridical.

These particular assumptions are likely to be incorrect, however, in ways that systematically underestimate the actual distance from viewers to ceilings versus floors. First, many ceilings are higher than 8 ft, particularly in public buildings, and almost none are lower than 8 ft. Second, most people spend much of their time inside rooms sitting rather than standing, which typically puts their eyes somewhat below the midway point between the floor and a standard ceiling, with any higher ceiling being considerably farther away than the floor. Third, people see many objects on underhead surfaces that are raised above the floor, such as tables and countertops, but see virtually no objects on overhead planes that are lowered below the ceiling. This fact further biases the statistics toward overhead planes

being farther away than underhead ones. Given all of these considerations, it is not unreasonable that, in the absence of more reliable depth information from stereopsis or motion, a veridical bias might exist toward seeing objects on overhead planes as being farther away than objects on underhead planes. Because we know of no ecological data on this question, however, we leave it as an open conjecture to be addressed by future research.

We note in passing that if the sizes of images of the same sphere in the room had been rendered accurately at different distances, they would have been smallest when they were closest to the horizon and largest when they were farthest from the horizon. This was not the case in our displays, however, because we purposely made them the same size (in pixels) in every image, to avoid confounds with the depth cue of relative size. We believe that this departure from realism does not detract from the conclusions of our study, because the results are, if anything, less extreme than they otherwise would have been had relative size been included as an additional depth cue. Its absence should, if anything, compress observers' depth judgments, rather than exaggerate them.

EXPERIMENT 2

Depth Perception Without Information About Distance to the Horizon

Experiment 1 demonstrated an overwhelming tendency for observers to base their depth judgments on DH rather than on HPP when the horizon was within the field of view between an overhead horizontal surface and an underhead horizontal surface. HPP might nevertheless be relevant in images in which the horizon is not explicitly present or in which DH is constant (e.g., Roelofs & Zee-man, 1957). To examine these possibilities, we conducted a second experiment.

The three conditions in Experiment 2 are shown in Figure 4. We eliminated all shading cues to surface slant and equalized all spatial cues except the vertical position of the object within the frame and the relative position of cast shadows that indicated whether the surface of attachment was below or above the sphere. In the *no-horizon* condition, there was no horizontal edge at all, the sphere and its cast shadow being shown on a flat, medium gray background (see Figure 4B). Here, HPP varies explicitly in the images, whereas DH does not, or at least not without an assumption about the location of the horizon relative to the picture boundaries. In the other two conditions, a horizontal edge was included that was always located at the center of the sphere (see Figures 4A and 4C). If this edge were perceived as the horizon of the supporting surface, the distance of the sphere relative to the horizon (DH) would be constant, whereas HPP would vary widely. The sphere in all three displays was lit more peripherally than in the previous experiment, providing a cast shadow that did not impart information about the slant of the surface on which it rested. As in Experiment 1, the ceiling plane images were reflections about a horizontal axis of the ground plane images, thereby changing the position of the cast shadow relative to the sphere.

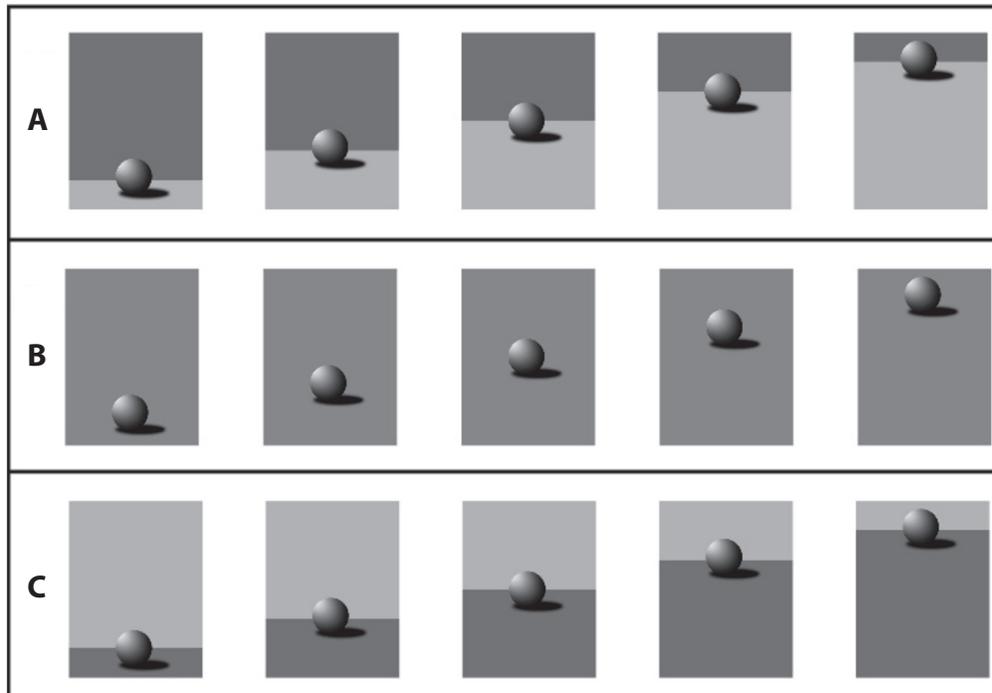


Figure 4. The three display conditions in Experiment 2. Height in the picture plane was systematically varied under conditions in which no horizon was present (B) or the horizon was always positioned at the center of the target sphere (A and C).

Method

Participants. All 24 participants were students at the University of California, Berkeley, who received partial course credit in their undergraduate psychology course. Their mean age was 19.3 years. As in Experiment 1, all had normal or corrected-to-normal vision, were naive as to the purpose and nature of the experiment, and gave informed consent in accord with the policies of the University of California, Berkeley, Committee for the Protection of Human Subjects.

Design. There were 120 paired comparisons, resulting from the permutations of the five vertical positions of the sphere, chosen two at a time, for each of the three horizon conditions (no horizon, darker bottom, and darker top) and for the two attachment conditions (floor and ceiling). Neither the different horizon conditions nor the different attachment conditions were compared with each other directly on any two-alternative forced choice trial.

Displays. All the images were created in Adobe Photoshop CS3 (www.adobe.com), and the monitor was the same as that in Experiment 1. Each image measured 360×480 pixels, and the positions were 80 pixels apart. The sphere measured 100 pixels in diameter and, as in Experiment 1, was colored gray, with shading to specify its curvature. The background in the no-horizon condition was homogeneous medium (0, 0, 58 in HSB coordinates) gray. The regions that defined the horizon in the other two conditions consisted of two flat grays, each of which was an equal distance from the medium gray from the no-horizon condition (0, 0, 41 and 0, 0, 75 in HSB coordinates). The only difference between these two conditions was the relative position of the grays in the images, which we call the *darker bottom* and the *darker top* conditions. Screen position was identical to that in Experiment 1, except that the images had 220 pixels separating them.

Procedure. The procedure was identical to that in Experiment 1.

Results and Discussion

The results for the floor conditions, in which the sphere's cast shadow was consistent with its resting on

a horizontal plane below eye level, were extremely clean and orderly (see Figure 5A). The average probability of choosing a given sphere as appearing farther away was a linear function of HPP, both when no horizon was present in the image [$F(1,23) = 857.0, p < .001$, for the linear component] and when a horizon was present but coincident with the center of the sphere [$F(1,23) = 445.0, p < .001$, also for the linear component]. As is shown in Figure 5A, the three floor conditions (the two horizon line conditions are shown as one line) were identical.

The results seem to imply the relevance of HPP when the spheres were on an underhead plane, because DH is technically undefined when no horizon is present (in the no-horizon conditions) and DH is constant when the sphere is at the same distance from the horizon (in the horizon conditions). Nevertheless, DH can also explain the results, provided that the observers judged the distance to the sphere using an unseen implicit horizon at a constant distance above the upper edge of the picture. This is especially likely in the no-horizon condition because of the lack of any horizontal edge in the image. An explanation in terms of DH is also possible in the constant horizon conditions, however, if the observers interpreted the horizontal edge in the image as the far edge of a finite surface (e.g., the far edge of a floor, tabletop, or countertop), whose actual horizon was at a constant height above the upper edge of the picture. In fact, people who live in urban environments seldom see the real horizon of a functionally infinite surface—such as the ocean or an enormous flat landscape of grass—so it would not be surprising if their default assumption were that a visible horizontal

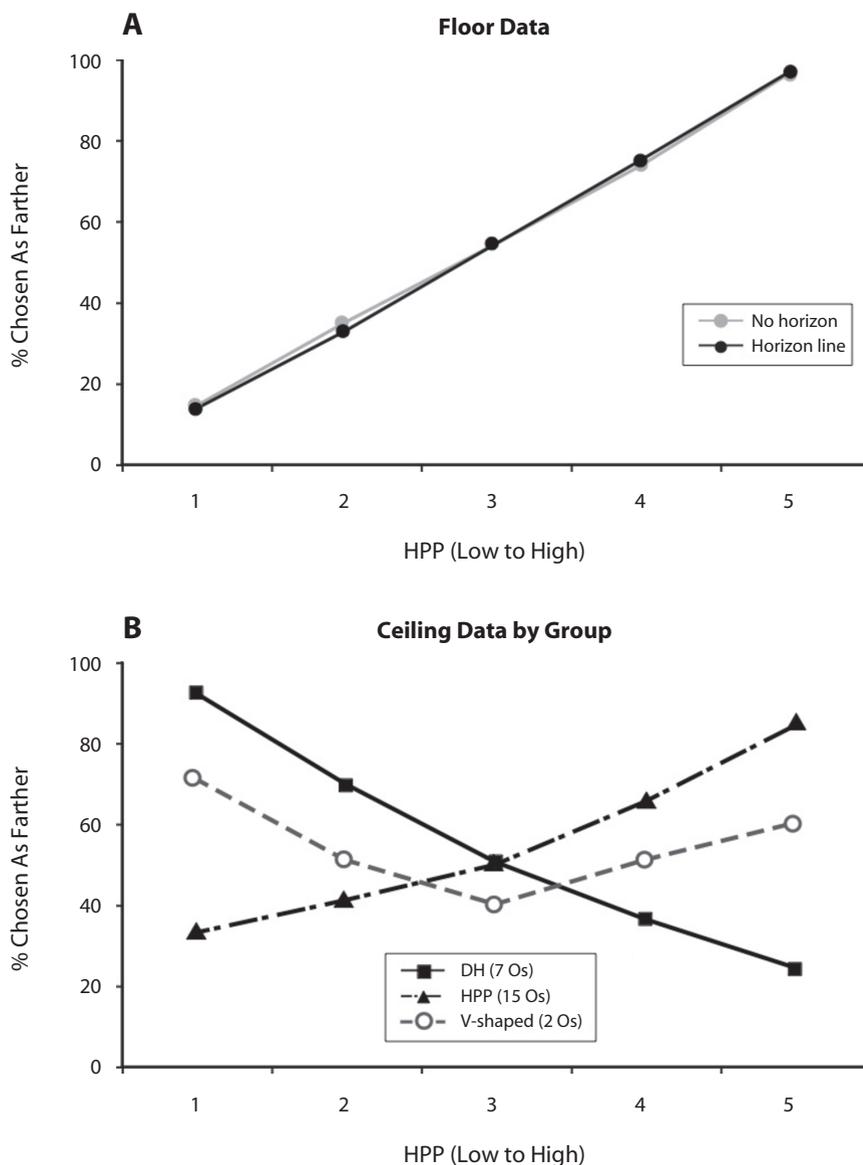


Figure 5. Results of Experiment 2. Average percentages of trials on which each image was chosen as depicting the sphere as farther away, when compared with other images from the same series in pairwise comparisons. (A) Results for the floor data for all participants. (B) Results for the ceiling data for three groups of participants who gave similar patterns of results. DH, distance to the horizon; HPP, height in the picture plane.

edge indicates the far edge of a finite surface, rather than the horizon of an infinite one. Another possibility is that they interpreted the horizontal edge as a reflectance edge on a single flat horizontal surface (e.g., a painted edge) whose actual horizon was at a constant height above the upper edge of the picture. The vertical position of the surface's horizon is indeterminate in all these cases, because no converging parallels or texture gradients are present to indicate the location of the true horizon. If observers take this unseen horizon to be at the same height in all cases, DH can also explain the results.

The ceiling data revealed a very different story, however, in that the observers showed three qualitatively dif-

ferent patterns of response (see Figure 5B). Seven observers made judgments that were clearly consistent with DH, producing a linear function that increased with increasing DH and decreased with increasing HPP. This pattern of response is fully consistent with the geometry implied by an unseen horizon at a constant distance below the bottom edge of the image.

Fifteen other observers made judgments that were in accord with HPP, producing a linear function that increased with increasing HPP and decreased with increasing DH. This pattern of response seems counterintuitive, given that the sphere's cast shadow clearly implies that it is attached to an overhead surface. It is as though these observers ig-

nored the shadows or did not know how to interpret them and responded to all the spheres as though they were supported by an underhead surface.

The remaining 2 observers showed a bizarre, V-shaped pattern of responses in which the apparent distance of the sphere first decreased with HPP and then increased. Perhaps they used DH to an unseen horizon below the bottom of the picture when the sphere was low in the frame and HPP when it was high in the frame. Interestingly, the observers in all three groups produced essentially identical data for the floor conditions (Figure 5A), in that there were no significant differences between their floor data.

Why might the patterns of response be so variable in the ceiling conditions of this experiment? People may be far less accustomed to judging the distance to objects above eye level than below eye level and, therefore, be less sensitive to the cues that specify it, especially when the cues are minimal. In the first experiment, the data for the sphere on the ceiling were very consistent across participants, perhaps because these images convey enough information from converging parallel edges to specify the horizon's vertical position in the picture. This may have made DH much more salient and unambiguous than in the second experiment. By eliminating (or at least greatly reducing) the influence of the horizon in the second experiment, depth judgments of objects attached to an overhead surface appear to be much less natural and, thus, more difficult to make. Similar results have been reported by Reichel and Todd (1990) in other experiments with inverted displays. They found that observers were biased toward seeing surfaces as slanting backward in depth rather than forward, despite texture and occlusion information dictating otherwise. The present pattern of results is thus consistent with other findings that show perception of ground planes to be superior to that of ceiling planes (e.g., Bian, Braunstein, & Andersen, 2005; McCarley & He, 2000).

SUMMARY AND CONCLUSION

The goal of the experiments reported here was to clarify the role of HPP versus DH as pictorial cues to depth. The results show that HPP clearly is *not* used when there is unambiguous information about the horizon and when surfaces both above and below eye level are present in the same image. Indeed, observers' depth judgments are almost perfectly consistent with the use of DH. When the target objects are located on an underhead plane, a proportional DH model accounts for virtually all of the variance, whereas when they are located on an overhead plane, some combination of proportional and absolute DH cues is required. For reasons that are not entirely clear, the observers consistently perceived objects on an overhead plane as farther away than correspondingly positioned ob-

jects on an underhead plane. When information about the position of the horizon is ambiguous, some people seem to use HPP to judge distance, even though it is geometrically incorrect, whereas others appear to use DH to an unseen horizon. In general, HPP appears to be relevant only to a limited number of observers viewing images in which information about the position of the horizon is absent, ambiguous, and/or unfamiliar. Additional research will be required to determine how general these findings are when the image transformations used to create the displays conform to naturally occurring eye and head movements (e.g., eye-head rotations and/or translations, rather than image croppings).

AUTHOR NOTE

This research was supported in part by NSF Grant BCS-0745820 and a Google gift to S.E.P. We thank Hye-Lim Jeon, Christine Nothelfer, Rosa Poggesi, Michael Vattuone, and other members of PalmerLab for their help in data collection. We also thank Hal Sedgwick and Jim Todd for their comments on an earlier draft of the manuscript. Correspondence concerning this article should be addressed to S. E. Palmer, Department of Psychology, University of California, Berkeley, CA 94720-1650 (e-mail: palmer@cogsci.berkeley.edu).

REFERENCES

- BIAN, Z., BRAUNSTEIN, M. L., & ANDERSEN, G. J. (2005). The ground dominance effect in the perception of 3-D layout. *Perception & Psychophysics*, *67*, 802-815.
- BLAKE, R., & SEKULER, R. (2006). *Perception* (5th ed.). New York: McGraw-Hill.
- DILDA, V., CREEM-REGEHR, S. H., & THOMPSON, W. B. (2005). Perceiving distances to targets on the floor and ceiling: A comparison of walking and matching measures [Abstract]. *Journal of Vision*, *5*(8), 196a.
- EPSTEIN, W. (1966). Perceived depth as a function of relative height under three background conditions. *Journal of Experimental Psychology*, *72*, 335-338.
- MAMASSIAN, P., KNILL, D. C., & KERSTEN, D. (1998). The perception of cast shadows. *Trends in Cognitive Sciences*, *2*, 288-295.
- MCCARLEY, J. S., & HE, Z. J. (2000). Asymmetry in 3-D perceptual organization: Ground-like surface superior to ceiling-like surface. *Perception & Psychophysics*, *62*, 540-549.
- OOI, T. L., WU, B., & HE, Z. J. (2001). Distance determined by the angular declination below the horizon. *Nature*, *414*, 197-200.
- PALMER, S. E. (1999). *Vision science: Photons to phenomenology*. Cambridge, MA: MIT Press.
- REICHEL, F. D., & TODD, J. T. (1990). Perceived depth inversion of smoothly curved surfaces due to image orientation. *Journal of Experimental Psychology: Human Perception & Performance*, *16*, 653-664.
- ROELOFS, C. O., & ZEEMAN, W. P. C. (1957). Apparent size and distance in binocular and monocular distance. *Ophthalmologica*, *133*, 183-204.
- SEDGWICK, H. A. (1986). Space perception. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance: Vol. 1. Sensory processes* (pp. 21.1-21.57). New York: Wiley.
- WOLFE, J. M., KLUENDER, K. R., LEVI, D. M., BARTOSHUK, L. M., HERZ, R. S., KLATZKY, R. L., ET AL. (2006). *Sensation and perception* (2nd ed.). Sunderland, MA: Sinauer.

(Manuscript received July 9, 2009;
revision accepted for publication September 25, 2009.)