

Distance perception across spatial discontinuities

J. C. MENG

Massachusetts Institute of Technology, Cambridge, Massachusetts

and

H. A. SEDGWICK

State University of New York, State College of Optometry, New York

We investigated the use of nested contact relations in perceiving the relative distance of locations on discontinuous surfaces. Observers viewed computer-generated displays under monocular static conditions and adjusted a marker to match the perceived distance of a cube. The marker and cube were raised above the ground by two different platforms separated by a gap. The relative heights and distances of the platforms were varied. We found the following: (1) When spatially discontinuous surfaces are coplanar, locations of objects resting on these surfaces appear to be compared directly, bypassing relations with the underlying ground plane. (2) Spatial displacement between the platforms produces a bias, in the direction of the displacement, in the perceived relative locations of objects resting on the platforms. This suggests that local spatial relations between objects and their platforms are only partially integrated with more global spatial relations between the discontinuous surfaces of the platforms.

Our previous studies in depth perception (Meng & Sedgwick, 2001) have supported a surface representation of 3-D layout. In such a representation, space is defined by surfaces, and distance is specified along those surfaces rather than through empty space (Alhazen, 1989, p. 155; Gibson, 1950, p. 6; Sinai, Ooi, & He, 1998; reviewed in Sedgwick, 2001). There is increasing evidence that visible surfaces play an important role in providing a framework for many other aspects of visual perception in a 3-D environment as well. Studies in stereopsis (Brookes & Stevens, 1989; Gillam & Sedgwick, 1996; Glennerster & McKee, 1999; Wilcox, Chodirker, & Bray, 2000), apparent motion (He & Nakayama, 1994), visual search (He & Nakayama, 1992), and visual attention (He & Nakayama, 1995) have all provided examples in which perception is influenced by or organized along visible surfaces (see review in Nakayama, He, & Shimojo, 1995).

In order to use a surface as reference, on top of which a 3-D representation is built, an object has to be related to that surface. There are many possible relations an object might have with respect to a surface. For example, an object can touch a surface, hover above a surface, be suspended from a surface, or move along a surface. We have

been particularly interested in studying the most direct and obvious relation between an object and a surface, which is the perceived contact. For example, a stationary object is seen as resting on the ground plane. For objects that are not directly resting on a common surface, we used the term *nested contact relations* (Meng & Sedgwick, 2001; Sedgwick, 1987, 1989) to describe a series of contact relations among multiple surfaces. In our previous experiments, observers were instructed to position a marker along a track on the ground to match the distance of a cube resting on a platform raised above the ground plane. Our results show that changing the perceived contact relation between the platform and the ground alone changes the perceived distance of the cube. Thus, the cube's perceived contact relation with the platform mediates its perceived location relative to the ground.

In those experiments, because the matching task was performed using a marker and track lying on the ground plane, the observers were forced to use the ground as the reference. In the experiments reported here, we explore situations where multiple reference surfaces are available. Our computer-simulated scenes show a cube resting on a platform on the left-hand side and a marker resting on a track on a platform on the right-hand side (Figures 1 and 2). The relative depth of the cube and the marker can be compared with respect to the ground plane, to the front surfaces of the two platforms, or to the top surfaces of the two platforms. Experiment 1 was designed to create situations where there is either a displacement between the top surfaces or the front surfaces of the two platforms. In theory, if the ground plane is always used as a global reference, performance should not be influenced by such displacements.

Portions of this work were presented at the 2000 meeting of the Association for Research in Vision and Ophthalmology, Fort Lauderdale, Florida. This research was supported in part by NSF Grant BCS-0001809 to H.A.S. We thank P. Kruger, N. Qian, and Q. Zaidi for helpful discussions and suggestions. We also thank G. J. Anderson, M. L. Braunstein, and J. Todd for their careful reviews and helpful comments. Correspondence should be addressed to H. A. Sedgwick, Department of Vision Sciences, State University of New York, State College of Optometry, 33 West 42nd St., New York, NY 10036 (e-mail: hsedgwick@sunyopt.edu).

EXPERIMENT 1

Method

We used the 3-D Studio Max software package (by Kinetix) to create virtual 3-D scenes. Objects were modeled in virtual 3-D space and rendered with photorealistic textures. During the experiment, pre-generated images were displayed on a 21-in. monitor one at a time. The pixel resolution used on the monitor was $1,024 \times 768$. A Windows NT program created with the PiXCL language (by Vysor Integration) controlled the sequencing and displaying of images and the gathering and storage of data.

Observers viewed the displays monocularly in the dark with their dominant eye. The observer's head was held steady by a chin and forehead rest. The perspective view simulated that of a 35-mm camera with a 50-mm ("normal") lens. The eye was aligned to be at the camera's point of view, which was 14.0 cm above the bottom of the monitor screen and 52.1 cm from the monitor screen. The total eye height above the floor was 118 cm. The whole scene subtended a visual angle of 39.5° horizontally and 30.1° vertically. A black hood hid the edges of the monitor, as well as the rest of the room.

There is no information in our displays to specify the absolute scale of the scene, nor do we measure the perceived scale of the scene. Because the task required in these experiments is a relative distance judgment, variation in perceived scale of the scene is not relevant and should not affect our results. The dimensions of all the scene parameters and the measured responses given in this paper are scaled according to an arbitrarily assumed camera height of 140 cm above the ground plane in the virtual space.

Each scene consisted of a textured ground plane extending to the horizon (Figures 1 and 2). Two large rectangular platforms rested parallel to each other, with a 100-cm separation in between. The left platform (1,600 cm long \times 150 cm wide \times 64 cm high) was mapped with a greenish marble texture, and the right one (1,600 cm long \times 150 cm wide \times various heights) with a pinkish graphite texture. Different textures were used to discourage the observers from matching locations on the platforms based on specific textural cues such as the projected sizes of texture elements. A small cube (40 \times 40 \times 40 cm) rested on top of the left platform, centered with respect to its left and right edges and 380 cm away from its front edge. A 20-cm-wide and 1,600-cm-long transparent track was centered on top of the right platform.

The relations between the two platforms were varied in two ways in Experiment 1. In Part 1, both platforms were at the same distance from the observer (Figure 1). The height of the right platform, however, varied between 16, 32, 48, and 64 cm (Figures 1c–1f). The corresponding configurations expressed as height ratios between the left and right platforms were 4:1, 4:2, 4:3, and 4:4. Two control conditions were added. In one, the right platform was absent and the track lay directly on the ground (height ratio 4:0, Figure 1b). In the other, both platforms were absent so that the cube and the track both lay directly on the ground (height ratio 0:0, Figure 1a).

In Part 2, the right platform and track were either 1 m closer to (forward conditions) or 1 m farther away from (backward conditions) the observer relative to the left platform (Figure 2). The forward and backward conditions were designed to be compared with the aligned conditions, where the left and right platforms were at the same distance from the observer. For each position of the right platform, three height ratios were used: 4:0, 4:2, and 4:4. For the aligned condition, the height ratios overlapped with those in Part 1, and therefore were not repeated.

At the beginning of each trial, a thin horizontal red marker (20 cm wide) appeared resting at either the near or the far end of the track on the right platform. The red marker could be moved along the track, changing size in perspective, by holding either the Up or Down arrow keys on the keyboard. The observers were instructed to use the arrow keys to move the red marker along the track on the right platform until it matched the perceived distance of the front edge of the cube on the left platform. We were careful not to sug-

gest any strategy in performing the task. To help the observers to conceptualize the task, however, they were encouraged to think of the display as representing a real 3-D scene in which they could walk along the ground between the platforms to a location where the cube was on their left and the red marker on their right. The experimenter also demonstrated the idea by using the space between tables in the laboratory. The observers were instructed to bracket their settings by moving the marker back and forth starting from locations that were obviously too far or too close; they were allowed to take as much time as they needed. When each setting was completed, the observer pressed the space bar on the keyboard to end the trial and to start the next one.

The 12 configurations illustrated in Figures 1 and 2 were each presented at eight distances along the ground. This was done to prevent the observers from remembering their previous settings. Distances were at 20-cm increments, with the cube's distance from the observer ranging from 11.01 to 12.41 m. We also varied the texture pattern on the ground and on the platforms from trial to trial so that the observers were not able to use texture landmarks to recognize their previous settings. For each configuration, observers made both an ascending and a descending adjustment, with the red marker starting at either the near or the far end of the track, respectively.

The combination of configuration (12) \times distance (8) \times starting position (two) gave rise to 192 trials, which were presented in a random order. Ten observers participated, including the 2 authors. All had 20/20 corrected vision with their dominant eye. Apart from the authors, the observers were naive as to the purpose of the experiment.

Results

The experimental trials for the 12 configurations illustrated in Figures 1 and 2 were combined and randomized together. The results for Parts 1 and 2, however, are presented and discussed separately.

Height variation of the right platform (Figure 1). The average settings, across eight distances, were calculated for each observer. On the basis of the averages for all the observers, the means, standard deviations, and 95% confidence intervals were calculated. These means and 95% confidence intervals across observers are plotted in Figure 3a. Since the results for ascending and descending trials showed no statistical difference, they were pooled. Each bar in the figure thus represents an average of ascending and descending trials over eight distances on the ground for all 10 observers (160 trials per bar). The dashed horizontal line marks the average simulated distance of the cube, which remained constant at 11.71 m across conditions. A three-way (configuration, distance, and starting position) within-subjects repeated measures analysis of variance (ANOVA) showed that the overall difference in mean matched distance among the six configurations was significant [$F(5,45) = 3.94, p < .005$]. Further LSD (least significant differences) post hoc analyses among the means showed that in Configurations 4:1 (12.29 m), 4:2 (12.20 m), and 4:3 (12.14 m), the matched distance was significantly greater than in Configurations 0:0 (11.64 m) and 4:4 (11.63 m). Configuration 4:0 (12.00 m) was roughly midway between the two groups and not statistically different from either group. This suggests that a continuous surface (0:0), or coplanar surfaces (4:4), provides a more accurate reference to compare the cube on the left platform to the marker on the right. When the cube and marker rested on surfaces with different heights (4:0,

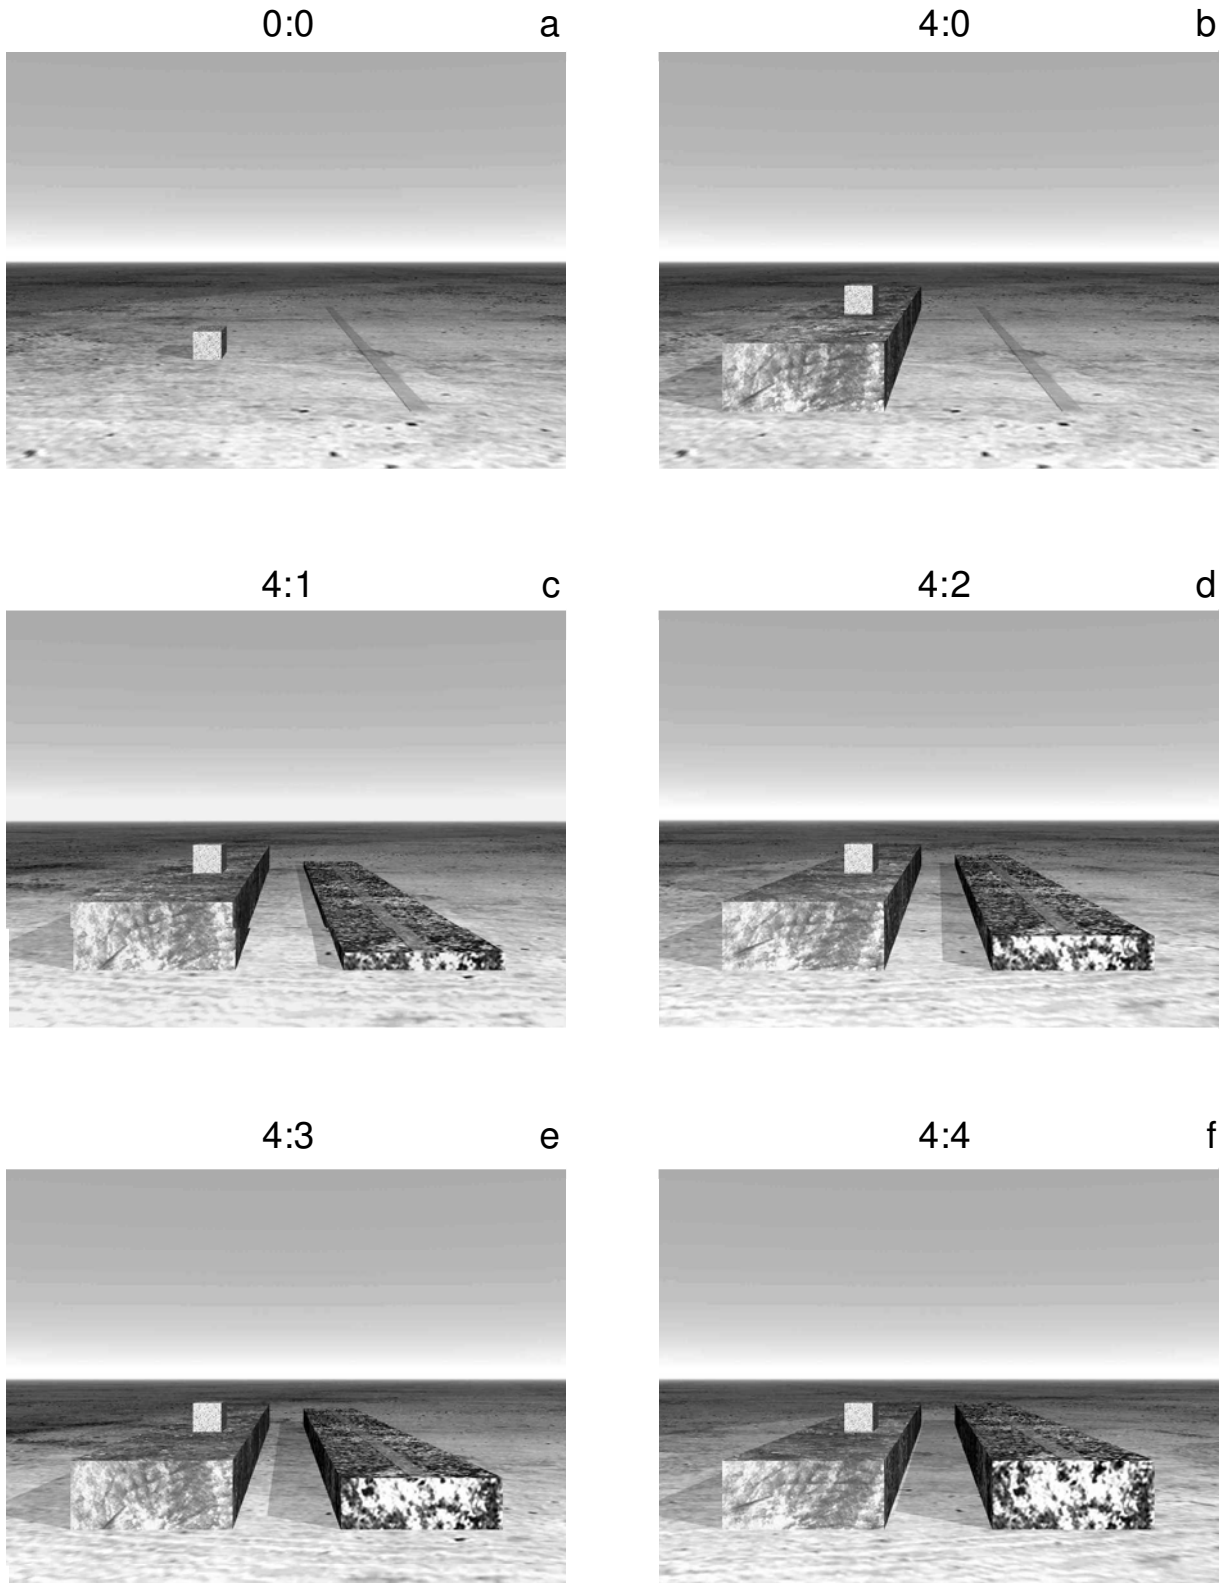


Figure 1. Spatial configurations used in Experiment 1, Part 1, where both the left and right platforms are at the same distance from the observers. (a) Both the cube and track rest directly on the ground plane. Since there is no platform present, this configuration is referred to as 0:0. (b) The cube rests on a platform 64 cm high and the track rests directly on the ground. The height ratio is referred to as 4:0. (c–f) The cube rests on the left platform, 64 cm high. The track rests on the right platform 16, 32, 48, and 64 cm high. The corresponding height ratios between the left and right platforms are 4:1, 4:2, 4:3, and 4:4. (Original scenes are in color.)

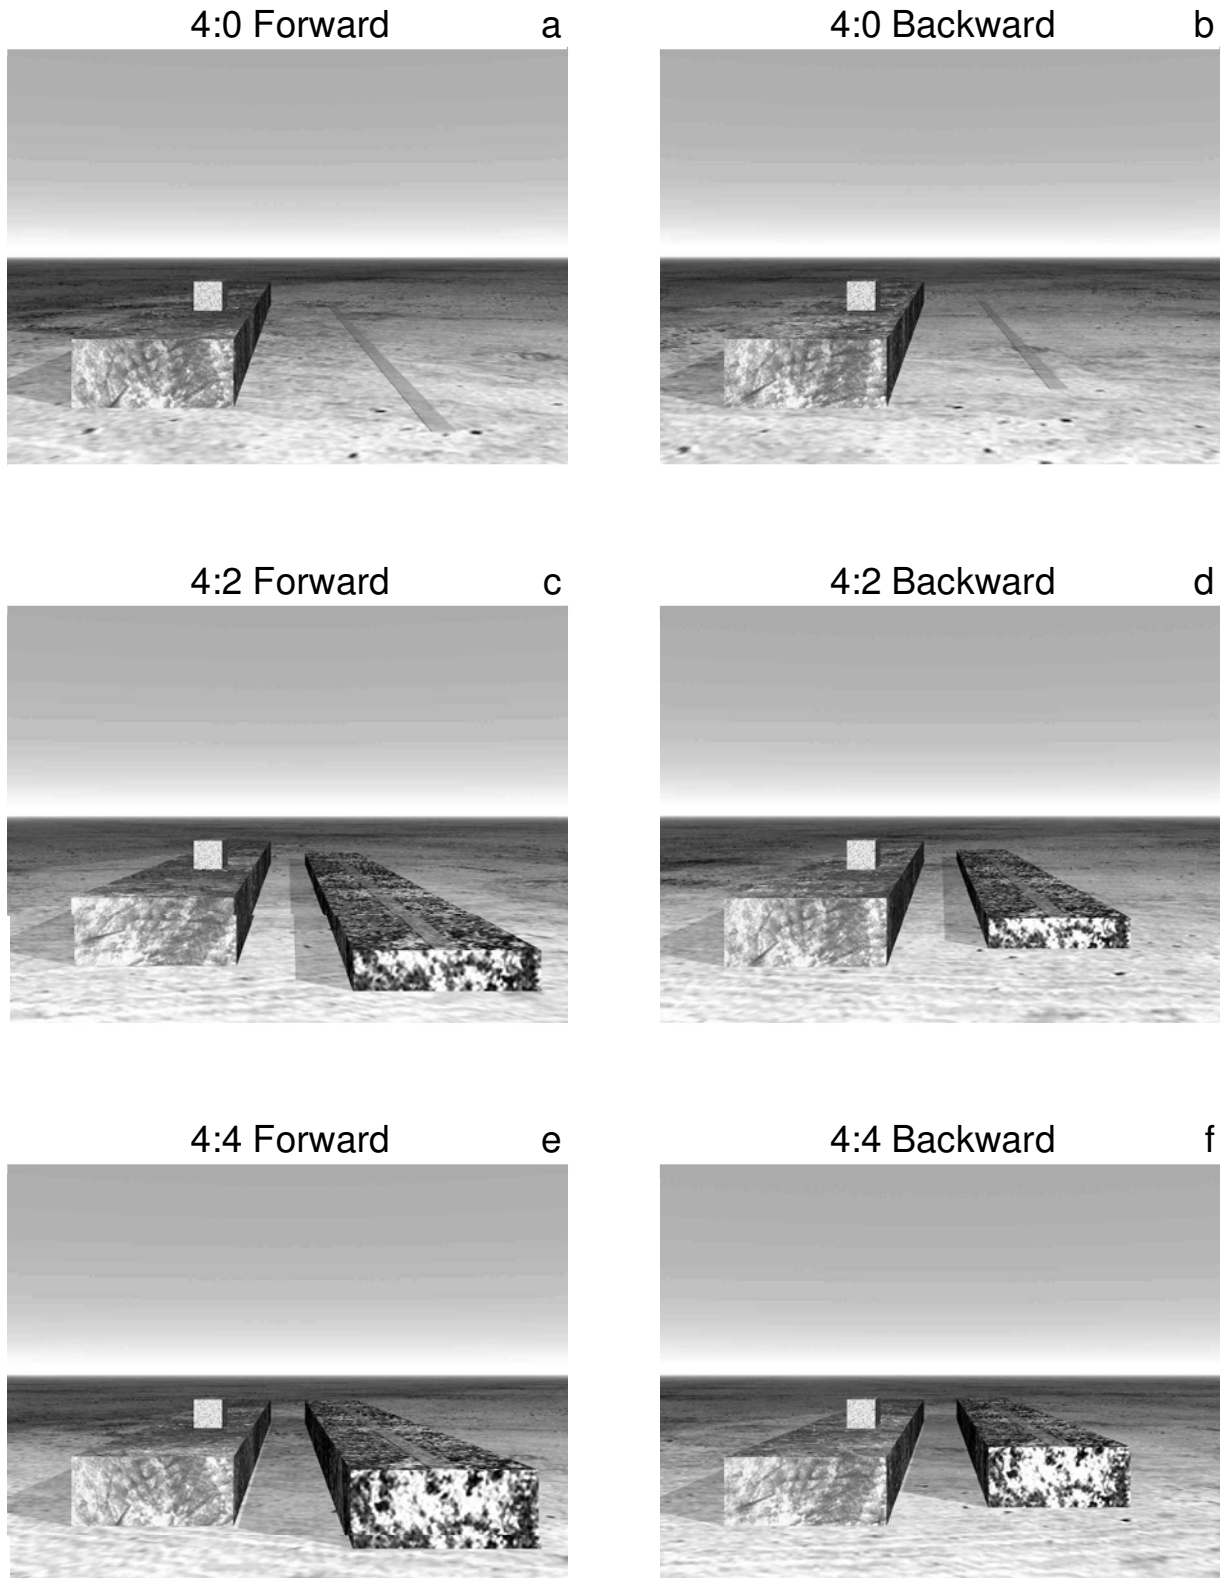


Figure 2. Spatial configurations used in Experiment 1, Part 2. The track and the right platform are either 1 m closer to (forward conditions) or 1 m farther from (backward conditions) the observers. The cube always rests on the left platform 64 cm high. (a–b) The track rests directly on the ground. The height ratio is referred to as 4:0. (c–d) The track rests on the right platform 32 cm high. The height ratio between the left and right platforms is 4:2. (e–f) The track rests on the right platform 64 cm high. The height ratio between the left and right platforms is 4:4. (Original scenes are in color.)

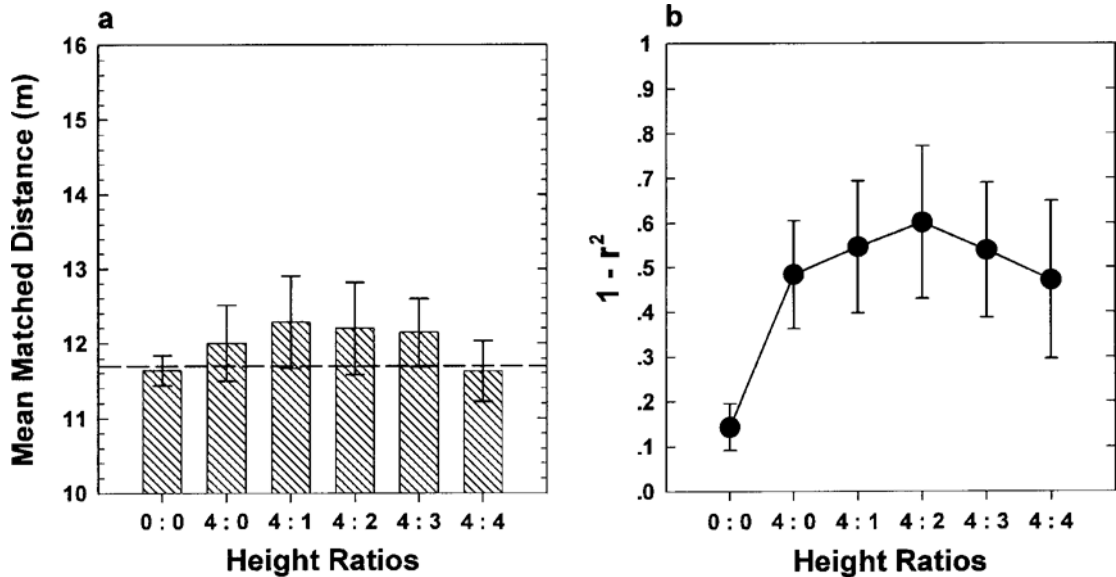


Figure 3. Results from Experiment 1, Part 1. (a) Means, with 95% confidence intervals across observers, of matched distance of the cube over eight distances along the ground for all 10 observers. The abscissa refers to the height ratios between the left and right platforms. The horizontal dashed line marks the average simulated distance of the cube at 11.71 m. (b) Means, with 95% confidence intervals across observers, of the residual variance, $1 - r^2$.

4:1, 4:2, and 4:3), performance was less accurate, with distance being overestimated. Additionally, in the experiments reported in this paper, the effect on mean matched distance of varying the distance of the configurations was always highly significant. This is to be expected and is not discussed further.

In order to assess the variability of our observers' performance, a linear regression line was fitted for each observer between the simulated and the matched distances. Squared correlation coefficients (r^2 s) were calculated to estimate the percentage of variability accounted for. Since a higher r^2 value stands for better precision in performance, we plotted the residual variance, $1 - r^2$, to indicate variability. Figure 3b shows the means, with 95% confidence intervals across observers, of the $1 - r^2$ values at each configuration over 10 observers. A one-way ANOVA (within-subjects repeated measure) showed that the overall means were significantly different among the six configurations [$F(5,45) = 8.89, p < .00001$]. The LSD post hoc analyses further showed that variability was significantly less in the 0:0 configuration ($1 - r^2 = .12$) than in the other configurations. In particular, variability in the 4:4 configuration ($1 - r^2 = .47$) was significantly higher than in the 0:0 configuration, even though the means (Figure 3a) of the two configurations were nearly identical. This suggests that a continuous surface is better than coplanar surfaces in reducing uncertainty and improving precision.

Variation in relative distal positions between the two platforms (Figure 2). The means, with 95% confidence intervals across observers, of matched distance of the cube for each position of the platform are plotted in Figure 4a. For each height ratio (4:0, 4:2, and 4:4), the data are arranged according to the relative positions between

the two platforms. *Forward* refers to the right platform being 1 m closer to the observers (Figures 2a, 2c, 2e). *Backward* refers to the right platform being 1 m farther away from the observers (Figures 2b, 2d, 2f). The data from the forward and backward conditions were compared with the data obtained in Part 1 at height ratios 4:0, 4:2, and 4:4. Since the two platforms were at the same distance from the observer in Part 1, those data are labeled as the *aligned* conditions (Figures 1b, 1d, 1f). In all the configurations tested here, the matched distances were clearly biased by the relative distal positions between the two platforms, showing underestimation when the right platform was closer and overestimation when the right platform was farther away. These displacements were roughly symmetrical around the means for the aligned condition. A four-way ANOVA confirmed that the effect of platform position was statistically significant [$F(2,18) = 11.154, p < .0007$]. LSD post hoc analyses further showed that this effect was statistically significant within each height configuration; that is, the means for the forward, aligned, and backward positions in each height configuration (4:0, 4:2, and 4:4) were significantly different from each other. In order to rule out the possibility of an averaging effect, we also looked at the individual data for each observer (not shown here) and found that 9 out of 10 observers showed the same pattern seen here in Figure 4a, although to various degrees. These results suggest that the local frameworks provided by the platforms and the track had a significant effect in biasing the integration of distance information across discontinuous spatial locations.

The overall matched distance varied significantly across height configurations according to the ANOVA [$F(2,18) = 4.398, p < .03$]. LSD post hoc analyses showed that Con-

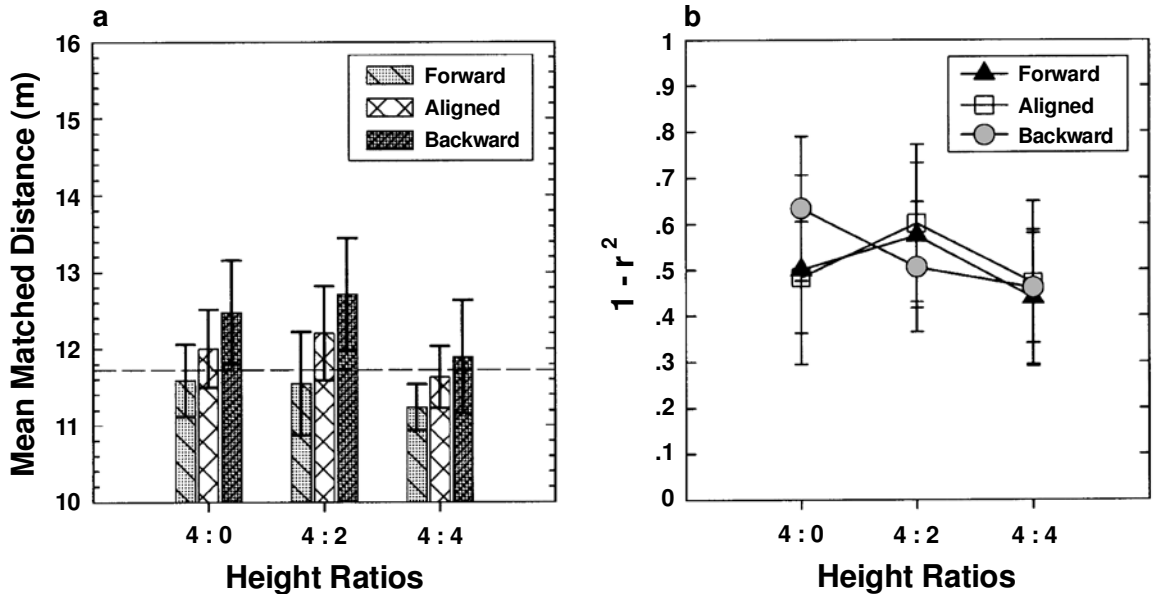


Figure 4. Results from Experiment 1, Part 2. (a) Means, with 95% confidence intervals across observers, of the matched distance of the cube over eight distances along the ground. At each height ratio, the data are arranged according to the relative position of the right platform. *Forward* refers to the right platform being 1 m closer to the observers, and *backward* to its being 1 m farther away. These two conditions are also compared with corresponding height configurations from Experiment 1, Part 1, where the two platforms are at the same distance from the observers (*aligned*). The horizontal dashed line marks the average simulated distance of the cube at 11.71 m. (b) Means, with 95% confidence intervals across observers, of the residual variance, $1 - r^2$.

Figurations 4:0 and 4:2 both differed significantly from Configuration 4:4, but not from each other. The ANOVA also showed a significant interaction between the height configurations and the distal positions of the right platform [$F(4,36) = 4.4001, p < .006$]. Figure 4a shows that the difference between the forward and backward conditions was the smallest in the 4:4 configuration, when the two platforms were of the same height (Figures 2e, 2f). Even displacing the track alone (Figures 2a, 2b), in the 4:0 configuration, produced a larger effect than in the 4:4 configuration. This supports the conclusion from Part 1, that coplanarity of the reference surfaces provides an advantage that allows the relative distance between objects to be compared more accurately across spatial discontinuities. The ANOVA also showed a significant four-way interaction [$F(28,252) = 1.94, p < .004$], for which we have no explanation.

The same statistical method used in Part 1 was employed to assess performance variability. Figure 4b shows the means, with 95% confidence intervals across observers, of the $1 - r^2$ values corresponding to all the conditions in Figure 4a. There was no statistical difference found among any of the conditions.

Discussion

These results show that observers are able to perceive spatial relations among objects across discontinuous surfaces. Precision, however, is compromised. Coplanar surfaces (when the top surfaces of the two platforms are of

the same height, or when the front surfaces are at the same distances from the observer) provide a significant advantage in relating the cube to the marker across the spatial gap. A continuous planar surface such as the ground plane, however, is an even more reliable reference.

The relative displacement between the two platforms in depth has a significant effect in biasing the results from the depth-matching task in the same direction. This suggests that the local framework provided by the platform biases the integration of distance information across discontinuous spatial locations.

We have also observed that in the coplanar configurations (4:4) the local framework bias is still present, but decreased. This suggests that, given a particular spatial configuration, observers may compare the relative distance between objects with respect to both reference surfaces rather than to only one or the other.

EXPERIMENT 2

Experiment 2 was designed to explore several issues raised by the results of Experiment 1. First, having found in Experiment 1 that performance was best in the 4:4 configuration, where the top surfaces of the two platforms were coplanar, we wished to test the generality of that finding; we did so by replicating the 4:4 configuration and adding a 2:2 configuration, in which the top surfaces were also coplanar but at the half the height. Second, we were somewhat surprised to find in Experiment 1 that there was a

strong effect of displacement even when the track rested directly on the ground in the 4:0 configuration. To further explore the generality of this effect in Experiment 2, we introduced a displacement into the 0:0 configuration, in which both the cube and the track rested on the ground. Furthermore, to make all the conditions more strictly comparable in terms of local frameworks, when the track or the cube rested on the ground, we inserted a very thin platform, like a carpet, under them.

Finally, in Experiment 1 we had used two configurations, 4:0 and 4:2, to investigate the interaction of different platform/track heights with forward/backward displacements of the platform and the track. These two configurations, however, differed both in amount of the height difference (4 or 2) and in whether the track rested on the ground (4:0) or on a platform above the ground (4:2). In Experiment 2 we replicated these two configurations while adding a third configuration (2:0) that allowed us to separate the effect of amount of height difference from the effect of having the platform resting on the ground.

Thus, in Experiment 2 we tested six height configurations: 0:0, 2:2, 4:4, 2:0, 4:0, and 4:2. Each of these was presented at our three displacements: forward, aligned, and backward.

Method

The viewing arrangement and method of adjustment for the task were the same as in Experiment 1. The scenes were similar to those in Experiment 1, but with different height configurations for the specific rationales discussed above. The left and right platforms had the same length, width, and texture as those in Experiment 1. In each pair of height configurations, the right platform was placed 1 m closer to (forward), at the same distance as (aligned), or 1 m farther away (backward) from the observers.

In the 0:0 configuration (Figures 5a, 5b), the two platforms shared the same very small height of 1 cm, which made them look more like a carpet resting on the ground than a platform. In the 2:2 configuration (Figures 5c, 5d), the two platforms shared the same height of 32 cm. This condition served as a comparison with the 4:4 configuration (Figures 5e, 5f), where the platforms shared the same height of 64 cm.

In the 2:0 and 4:0 configurations (Figures 6a, 6b, 6c, 6d), the left platform was 32 and 64 cm in height, respectively. The right platform was 1 cm in height (carpet). In the 4:2 conditions (Figures 6e, 6f), the heights of the left and right platforms were 64 and 32 cm, respectively.

The observer's task remained the same, which was to adjust the marker along the track on the right platform to match the perceived distance of the cube's front edge on the left platform. The six configurations were presented at each of eight distances along the ground, ranging from 11.01 to 12.41 m in 20-cm increments. Texture patterns were changed from trial to trial. For each scene, the observer made an ascending and a descending adjustment with the marker starting at either the near or the far end of the track, respectively.

The combination of height configuration (six) \times right platform position (three) \times distance (eight) \times starting position (two) gave rise to 288 trials, which were presented in a random order. A different group of 10 naive observers participated. All had 20/20 corrected vision with their dominant eye.

Results

The means, with 95% confidence intervals across observers, of matched distance of the cube are plotted in Figure 7a. For each configuration, the data are arranged ac-

cording to the relative distal positions between the two platforms. *Forward* refers to the right platform being 1 m closer to the observer (Figures 5a, 5c, 5e, Figures 6a, 6c, 6e). *Aligned* refers to the two platforms being at the same distance from the observer (analogous to Figure 1). *Backward* refers to the right platform being 1 m farther away from the observer (Figures 5b, 5d, 5f, Figures 6b, 6d, 6f). Again, the relative displacement in depth between the two platforms biases the matched distance of the cube significantly [three-way ANOVA, $F(2,18) = 38.74, p < .0001$].¹ LSD post hoc analyses further showed that this effect was statistically significant within each height configuration; that is, the mean matched distances for the forward and backward positions in each height configuration were significantly different from each other.

The mean matched distances across the six height configurations were significantly different [$F(5,45) = 20.616, p < .00001$]. LSD post hoc analyses showed that in Configurations 2:0, 4:0, and 4:2, observers were significantly overestimating in comparison with Configurations 0:0, 2:2, and 4:4.

The ANOVA also showed a significant interaction between the height configurations and the distal positions of the right platform [$F(10,90) = 5.515, p < .00001$]. Figure 7a indicates that the amount of the framework effect is the smallest in the 0:0 configuration, where the cube and track rested on the carpets, 1 cm above the ground surface. The order of the magnitude of the framework effect was 0:0 < 2:2 < 4:4 < 2:0 < 4:2 < 4:0. We again looked at each individual's performance and found that all 10 observers exhibited the framework effect. Two observers showed no significant effect in the 0:0, 2:2, and 4:4 configurations, but did show a significant effect in the 2:0, 4:2, and 4:0 configurations. The rest of the observers also showed much less of an effect with the cube and marker resting either on a continuous surface (0:0) or on coplanar surfaces (2:2 and 4:4).

The same statistical method used previously was employed to assess performance variability. Figure 7b shows the means, with 95% confidence intervals across observers, of the $1 - r^2$ values corresponding to all the conditions in Figure 7a. A one-way ANOVA showed that overall there were significant differences in variability across configurations [$F(5,45) = 27.98, p < .000001$]. The LSD post hoc analyses further showed that the performance variability was significantly less in the 0:0 configuration than in the coplanar configurations (2:2 and 4:4), which in turn were significantly less variable than the remaining configurations (2:0, 4:0, and 4:2). All 10 observers' individual data also demonstrate such a trend, although to various degrees. Figure 7b also shows that within each height configuration the aligned condition was slightly less variable than the forward and backward conditions, except in the 4:0 configuration.

Discussion

With a different set of 10 naive observers, the local framework effect introduced by the displacement in depth between the two platforms was observed again in the same

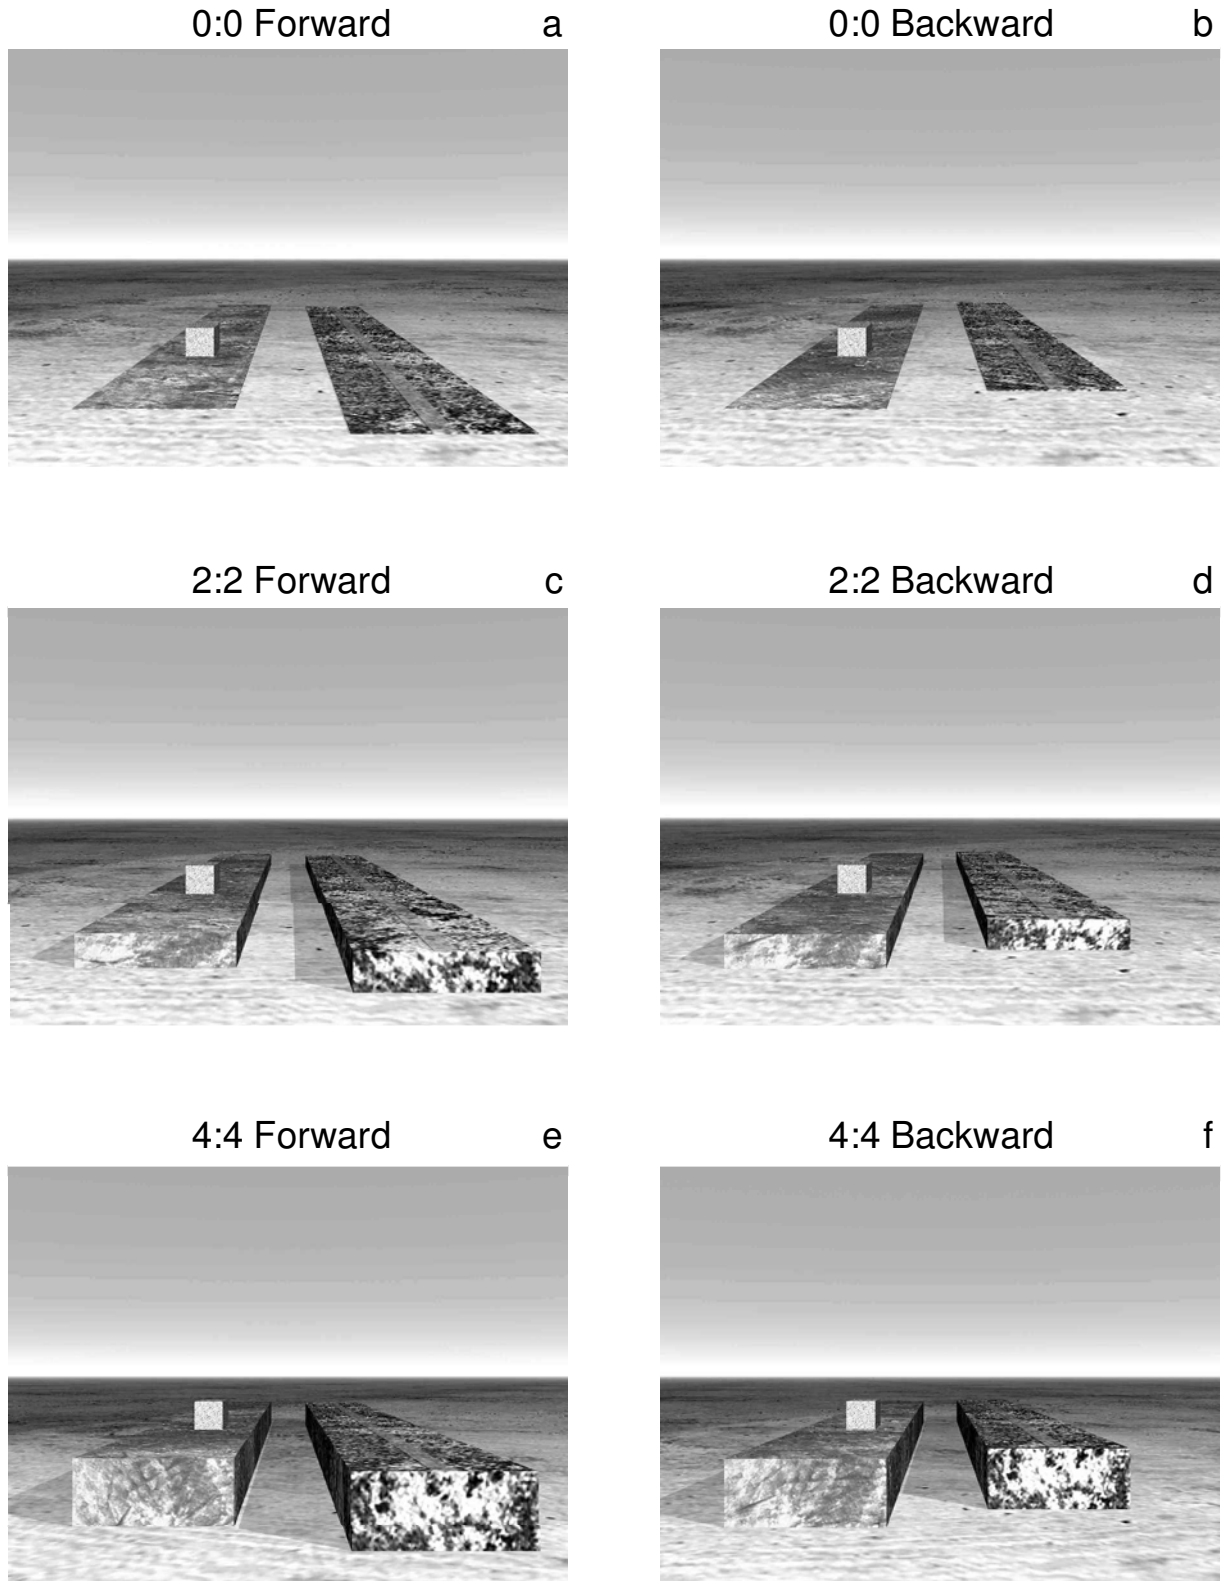


Figure 5. Spatial configurations used in Experiment 2, where the left and right platforms are of the same height. *Forward* and *backward* refer to the conditions where the right platform is 1 m closer to or 1 m farther away from the observers. (a–b) Both the left and right platforms are 1 cm high. This makes them look more like carpets than platforms. The height ratio is thus expressed as 0:0. (c–d) Both the left and right platforms are 32 cm high. Height ratio is expressed as 2:2. (e–f) Both the left and right platforms are 64 cm high. Height ratio is expressed as 4:4. *Aligned* conditions are not shown. (Original scenes are in color.)

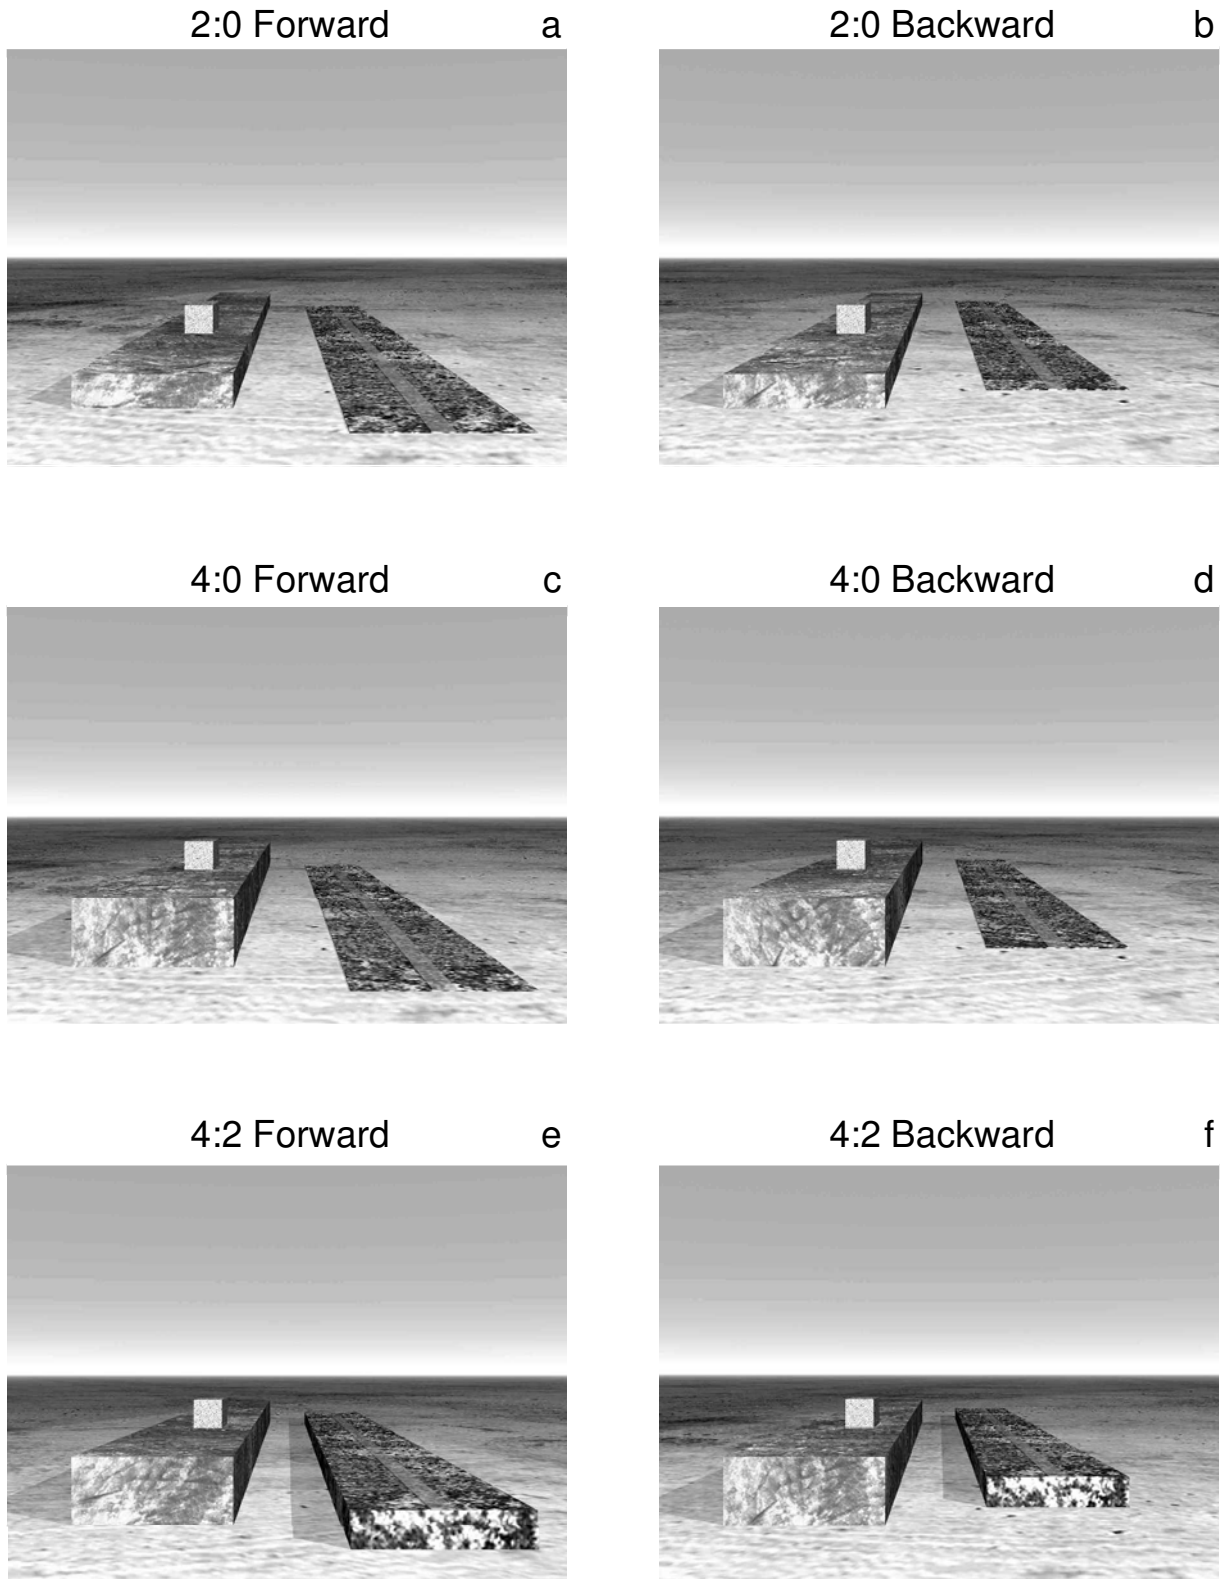


Figure 6. Spatial configurations used in Experiment 2, where the left and right platforms are of different heights. *Forward* and *backward* refer to the conditions where the right platform is 1 m closer to or farther away from the observers. (a–b) The left platform is 32 cm high, and the right one is 1 cm high. The ratio between the two platforms is expressed as 2:0. (c–d) The left platform is 64 cm high, and the right one is 1 cm high. The ratio between the two platforms is expressed as 4:0. (e–f) The left platform is 64 cm high, and the right one is 32 cm high. The ratio between the two platforms is expressed as 4:2. *Aligned* conditions are not shown. (Original scenes are in color.)

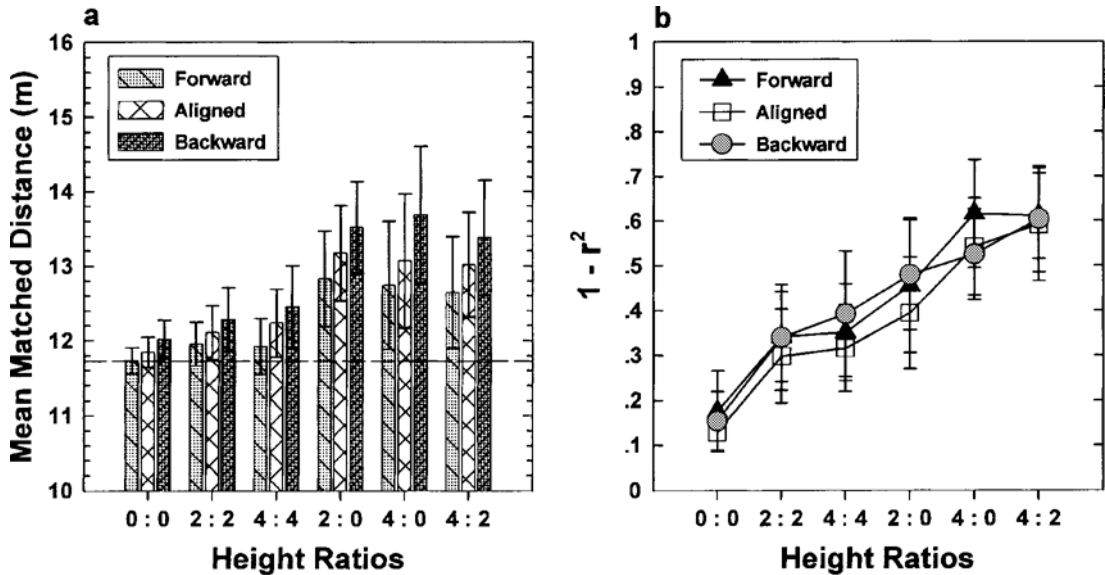


Figure 7. Results from Experiment 2. (a) Means, with 95% confidence intervals across observers, of the matched distance of the cube over eight distances along the ground. At each height ratio, the data are arranged according to the relative position of the right platform. *Forward* and *backward* refer to the right platform being 1 m closer to or farther away from the observers. *Aligned* refers to the condition where the two platforms are at the same distance from the observers. The horizontal dashed line marks the average simulated distance of the cube at 11.71 m. (b) Means, with 95% confidence intervals across observers, of the residual variance, $1 - r^2$.

direction. Such a bias manifested itself even in the 0:0 configuration, where the displacement was defined mainly by texture discontinuities outlined by the carpets resting only 1 cm above the ground plane. The continuity of the ground plane, however, reduced such a bias significantly. The coplanarity (2:2 and 4:4 configurations) of the top surfaces of the platforms also helped to reduce the local framework bias. This supports our findings in Experiment 1 that relative distance between objects resting on coplanar surfaces can be compared directly across spatial gaps.

In our previous studies (Meng & Sedgwick, 2001), we found that the distance between relevant surfaces influenced observers' performance variability. This finding was also observed by comparing the 2:0 and 4:0 configurations. The $1 - r^2$ value in the 2:0 configuration was significantly lower than that in the 4:0 configuration in each of the forward, aligned, and backward conditions tested.

The 4:4, 4:0, and 4:2 configurations were used in both Experiments 1 and 2. The only difference was in the 4:0 configuration. In Experiment 1, the track lay directly on the ground (Figures 2a, 2b). In Experiment 2, the track lay on top of a carpet, and the carpet in turn lay on the ground (Figures 6c, 6d). Figure 8a shows the overall means of matched distance for these three configurations in each experiment. The horizontal dashed line marks the simulated distance of the cube at 11.71 m. Three four-way between-groups ANOVAs, one for each of the three height configurations, showed that observers from Experiment 2 were significantly overestimating relative to observers in Experiment 1 [4:4, $F(1,18) = 4.67, p < .04$; 4:0, $F(1,18) =$

7.56, $p < .01$; and 4:2, $F(1,18) = 5.21, p < .03$]. None of these three ANOVAs, however, showed a significant interaction between the experiment factor (Experiment 1 vs. Experiment 2) and the distal positions of the right platform. This means that the biasing effects found in Experiment 1 are similar to those found in Experiment 2. We have previously observed such variation in overall estimation between experiments (Meng & Sedgwick, 2001), and we have no explanation for it.² It may be related to random variation in the population, given the relatively small sample sizes ($N = 10$ observers) that we have used.

For each of these three configurations (4:4, 4:0, and 4:2) in Experiments 1 and 2, the mean $1 - r^2$ values, indicating the variability across observers, are plotted in Figure 8b. Three between-groups ANOVAs showed no statistical differences between Experiments 1 and 2 for any of these three configurations.

GENERAL DISCUSSION

In our previous work (Meng & Sedgwick, 2001) we investigated a variety of situations in which an object rests on top of a platform, which in turn rests on the ground plane. We showed that the perceived location of the object relative to the ground plane can be effectively mediated by the platform on which it rests, but the precision of this perception deteriorates as the height of the platform increases.

In the experiments reported here, we have examined a more complex situation, in which observers attempt to relate the location of an object resting on one platform to

that of a marker resting on another platform; both platforms are resting on the ground with a gap, or spatial discontinuity, between their top surfaces. We found that both the accuracy and the precision with which observers can do this task depend on the spatial relations between the relevant surfaces of the two platforms. Specifically, performance is best when both the front surfaces and the top surfaces of the two platforms are aligned, or coplanar.

How can we account for the effects that we found? In our previous work (Meng & Sedgwick, 2001), we considered two different ways in which observers might link the cube resting on top of the platform to the marker lying on the ground. One way involves what we could call "alignment matching." Briefly, this requires the observer, using information about the height and configuration of the platform, to find the location on the ground directly under the cube. The marker can then be aligned with this ground location. The second way involves what we could call "depth matching." This requires the observer to match the perceived distance of the marker to the sum of two distances: the distance from the observer to the front of the platform and the distance from the front of the platform to the cube. An analysis of our observers' individual data suggested that individual observers might use both types of matching but that some observers rely more heavily on alignment matching and others rely more heavily on depth matching.

One indicator of which type of matching an observer was using came from the effect of increasing the height above the ground of the platform on which the cube rested. As the height of the platform above the ground increases,

the line of sight to the platform becomes more shallow, leading to increasing projective compression of the depth dimension of the platform. As Loomis and Philbeck (1999) have shown, the amount of this compression is directly related to the underestimation of distance. Thus, observers using depth matching would be expected to show a pattern of increasing underestimation of distance linked to increasing height of the platform. In a detailed analysis of the individual data, several of our observers did show this pattern to a statistically significant degree. On the other hand, other observers did not show this pattern, as would be expected if they were relying instead on alignment matching.

We can generalize both hypothetical types of matching to the two-platform configurations used in this paper. In alignment matching the ground locations of the cube and marker are each determined, presumably relatively independently, and the marker is adjusted until its ground location is aligned with that of the cube. In depth matching, the sum obtained from adding the perceived distance from the observer to the front of the right platform and the perceived distance from the front of that platform to the marker is adjusted to be equal to the sum obtained from adding the perceived distance from the observer to the left platform and the perceived distance from the front of that platform to the cube. We can now ask how well either or both of these types of matching might be able to account for the effects that we found.

First we will consider the possible use of alignment matching. This offers us a way of accounting for the improvements in performance that we found in both experiments when the top surfaces of the platforms were copla-

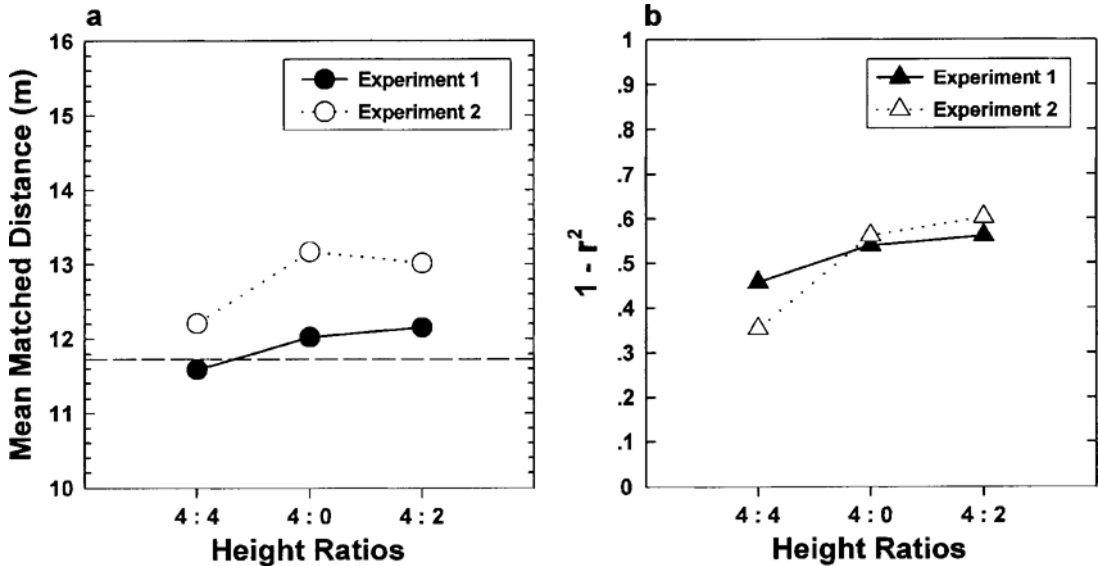


Figure 8. Comparison between two different groups of 10 observers from Experiments 1 and 2 at the corresponding spatial configurations. (a) Mean matched distance of the cube over eight distances along the ground for all 10 observers in each experiment. The horizontal dashed line marks the average simulated distance of the cube at 11.71 m. (b) Mean residual variance, $1 - r^2$, for the 10 observers in each experiment.

nar. We can do this by assuming that, when observers perceive the tops of the two platforms to be coplanar, they align the front edge of the marker directly with the front edge of the cube, rather than determining and aligning their respective ground locations. This shortcut would have the effect of increasing the precision of the observer's settings because it would largely remove the variability associated with linking a location on the top surface of each platform with the ground location directly beneath it. On the other hand, when the tops of the platforms are non-coplanar, observers would be unable to use this shortcut and so would determine and align the ground locations, thus increasing the variability of performance, as we found.

Because alignment matching makes no use of the positions of the fronts of the platforms, however, it should not be affected by the displacements between the two platforms in the forward and backward conditions. Thus alignment matching does not predict the bias effects produced by the relative displacements between the two platforms, and so cannot entirely account for our results.

Next we will consider the possible use of depth matching. In the first place it can be observed that a shortcut is possible in this type of matching. Rather than using the perceived distances from the observer to the left and right platforms, it would be sufficient to use the relative distance between the left and right platforms (D_{L-R}). Then the distance of the marker from the front of the right platform (P_R) could simply be adjusted to equal the sum of D_{L-R} and the distance from the front of the left platform to the cube (P_L). That is, the marker would be adjusted so that

$$P_R = P_L + D_{L-R}. \quad (1)$$

How might depth matching account for the biasing effects of the displacement? We see no reason to believe that the relative distance between the two platforms (D_{L-R}) is systematically misperceived. Nor do we see how to account for the bias by a systematic misperception of the distance from the front of either platform to the object (cube or marker) resting on it (P_R and P_L).³ Instead, we hypothesize that the bias arises from an incomplete integration of the local depth information, P_R and P_L , with the more global distance information, D_{L-R} , relating the positions of the two platforms.

Expressed in terms of Equation 1, the bias arises from a shortfall in the summing operation. We can express this by adding another term, k , where k represents some value less than 1, to the following equation:⁴

$$P_R = P_L + k * D_{L-R}, k < 1. \quad (2)$$

This expresses somewhat more precisely what we have referred to earlier in this paper as a *local framework effect*, which we take to mean that the relations between objects and their local frameworks are not completely integrated perceptually with the global relations between the frameworks. If k equals 1, then Equation 2 reduces to

Equation 1, and there is no bias produced by the local framework. If k equals 0, the setting of the marker is entirely determined by the local frameworks (i.e., the platforms), and the relative displacement between the platforms is not taken into account at all; we would then expect the amount of bias in the positioning of the marker to equal the relative displacement between the platforms. For values of k between 1 and 0, there is a partial integration of the global relative displacement information with local depth information so that the bias is correspondingly reduced. In our experiments, the average bias for each configuration is always less than the relative displacement between the platforms, thus implying values of k between 0 and 1.

Such local framework effects have been found in research on other aspects of visual perception. For example, Rock and Ebenholtz (1959) showed similar effects in comparing the length of lines surrounded by rectangular frames. A difference in the sizes of the rectangular frames produced a corresponding, although smaller, difference in the lengths of the lines that were perceived as equal. Although the rectangular frames were perceived as different in size, this perceived difference between frames was not fully integrated with the information relating each frame locally to the line that it contained. More recently, Whitney and Cavanagh (2000) have shown that motion between two horizontally aligned bars produces a bias, in the same direction as the motion, in the perceived relative positions of the bars.

In our experiment, this incomplete integration of local and global distance information would have the effect that in the forward configurations, in which the right platform (with the marker) is positioned nearer to the observer than is the left platform (with the cube), the perceived matching distance would be underestimated. And in the backward configurations, in which the right platform is positioned farther away than the left platform, the perceived matching distance would be overestimated. This is just the pattern of bias that we found in both experiments.

What depth matching cannot readily account for is the varying amount of the bias that we found for different height configurations. Specifically, the bias is significantly less in those configurations in which the two platforms are the same height (0:0, 2:2, and 4:4). This consideration leads us to hypothesize that both depth matching and alignment matching were being used by our observers, but that one or the other received more emphasis, depending on the configuration.

Thus, when the top surfaces of the two platforms are coplanar, more emphasis is given to alignment matching, which can produce a relatively precise match in this situation. Even here, however, there cannot be an exclusive use of alignment matching because a significant, although much smaller, bias effect is still present when there is a relative displacement between the two platforms.

Conversely, when the top surfaces of the two platforms are not coplanar, more emphasis is given to depth match-

ing, thus producing larger bias effects. There is also some evidence here, although perhaps less compelling, that some use is still made of alignment matching. Specifically, our finding that, for unequal platform heights, matching is less precise for greater inequalities (4:0 vs. 2:0) is at least consistent with the suggestion, supported by our earlier work, that the determination of the ground location of an object resting on a platform becomes less precise as the height of the platform above the ground increases.

Finally, we can hypothesize that in those situations in which the top of the platform is perceptually coplanar with the ground, like a carpet, there is more effective integration of depth information relating the object (cube or marker) to the platform with the relative distance information relating the two platforms, or carpets. This seems plausible here because all of the distance relations are coplanar. Thus in these situations we would expect the depth matching to be somewhat less biased by relative displacements of the platforms. This may account for our finding that the performance is better when both platforms are resting on the ground (0:0 configuration) than when the two platforms are coplanar but above the ground (2:2 and 4:4 configurations).

In analyzing the results of Experiments 1 and 2, as noted, we have looked at the individual data, as well as the averaged group data, and have consistently found that the individual observers show results whose patterns are at least qualitatively similar to the pattern of the group averages. This suggests to us that our hypothesis that both types of matching are being used, but with varying emphasis depending on the particular configuration of surfaces, applies to the individual observers and is not merely a description of overall group behavior such as might be produced by averaging together some observers performing in one way with other observers performing in another.

In summary, we have extended our earlier investigations of distance perception mediated by nested contact relations (Meng & Sedgwick, 2001) to more complex spatial layouts. We have shown that observers are able to make use of pictorially available information to match the distances of objects across spatial discontinuities arising from gaps between the platforms on which they rest. As in our earlier work, we have found that although observers make use of information linking objects to the ground plane, and hence to each other, they do not always use such information to its full potential. There are costs associated with this increased complexity as well as benefits associated with spatial configurations that simplify this complexity. In particular, we have identified strong local framework effects that arise with increased complexity in the scene; observers appear to relate objects first to their local frameworks and are not fully effective in integrating these local relations with the global relations between frameworks. Particular spatial configurations such as coplanarity, however, can create higher order configurations that allow observers to directly relate objects across local frameworks, thus reducing the local framework effects. Observers' performance reflects some flexibility in emphasizing one, or another, or a combination of possible reference surfaces.

REFERENCES

- ALHAZEN, I. (1989). Book of optics. In *The optics of Ibn-Haytham* (A. I. Sabra, Trans.). London: Warburg Institute, University of London. (Alhazen lived approximately 965–1040 C.E.; when he wrote this work is not known.)
- BROOKES, A., & STEVENS, K. A. (1989). Binocular depth from surfaces versus volumes. *Journal of Experimental Psychology: Human Perception & Performance*, **15**, 479–484.
- GIBSON, J. J. (1950). *The perception of the visual world*. Boston: Houghton Mifflin.
- GILLAM, B., & SEDGWICK, H. A. (1996). The interaction of stereopsis and perspective in the perception of depth. *Perception*, **25** (Suppl.), 70.
- GLENNERSTER, A., & MCKEE, S. P. (1999). Bias and sensitivity of stereo judgements in the presence of a slanted reference plane. *Vision Research*, **39**, 3057–3069.
- HE, Z. J., & NAKAYAMA, K. (1992). Surfaces versus features in visual search. *Nature*, **359**, 231–233.
- HE, Z. J., & NAKAYAMA, K. (1994). Apparent motion determined by surface layout not by disparity or three-dimensional distance. *Nature*, **367**, 173–175.
- HE, Z. J., & NAKAYAMA, K. (1995). Visual attention to surfaces in 3-D space. *Proceedings of the National Academy of Sciences*, **92**, 11155–11159.
- LOOMIS, J. M., & PHILBECK, J. W. (1999). Is the anisotropy of perceived 3-D shape invariant across scale? *Perception & Psychophysics*, **61**, 397–402.
- MENG, J. C., & SEDGWICK, H. A. (2001). Distance perception mediated through nested contact relations among surfaces. *Perception & Psychophysics*, **63**, 1–15.
- NAKAYAMA, K., HE, Z. J., & SHIMOJO, S. (1995). Visual surface representation: A critical link between lower-level and higher-level vision. In S. M. Kosslyn & D. N. Osherson (Eds.), *Visual cognition: An invitation to cognitive science* (2nd ed., Vol. 2, pp. 1–71). Cambridge, MA: MIT Press.
- ROCK, I., & EBENHOLTZ, S. (1959). The relational determination of perceived size. *Psychological Review*, **66**, 387–401.
- SAUER, C. W., BRAUNSTEIN, M. L., & ANDERSON, G. J. (1999). Judging the depth separation of objects in uniform connected regions in 3D scenes. *Investigative Ophthalmology & Vision Science*, **40**, S415.
- SEDGWICK, H. A. (1987). Layout2: A production system modeling visual perspective information. In *Proceedings of the First International Conference on Computer Vision* (pp. 662–666). Washington, DC: IEEE Computer Society Press.
- SEDGWICK, H. A. (1989). Combining multiple forms of visual information to specify contact relations in spatial layout. In P. S. Schenker (Ed.), *Sensor fusion II: Human and machine strategies* (SPIE Proceedings, Vol. 1, 198, pp. 447–458). Bellingham, WA: SPIE.
- SEDGWICK, H. A. (2001). Visual space perception. In E. B. Goldstein (Ed.), *Blackwell handbook of perception* (pp. 128–167). Oxford: Blackwell.
- SINAL, M. J., OOI, T. L., & HE, Z. J. (1998). Terrain influences the accurate judgement of distance. *Nature*, **395**, 497–500.
- WHITNEY, D., & CAVANAGH, P. (2000). Motion distorts visual space: Shifting the perceived position of remote stationary objects. *Investigative Ophthalmology & Vision Science*, **41**, S741.
- WILCOX, L. M., CHODIRKER, L., & BRAY, B. (2000). Stereoscopic surface interpolation. *Investigative Ophthalmology & Vision Science*, **41**, S735.

NOTES

1. In this ANOVA, ascending and descending trials were averaged together and not included as a factor in the analysis. This factor was included in our other analyses and did not approach significance in any of them.
2. Sauer, Braunstein, and Anderson (1999) have shown in computer simulations that perceived depth separations of two vertical poles standing on the ground are increased (by about 10%) when the poles are surrounded by a rectangular area in proper perspective. By analogy to their situation, it might be expected that depth intervals along the track in our Configuration 4:0 would appear somewhat longer in Experiment 2 than in Experiment 1 because of the presence of the rectangular carpet surrounding the track in Experiment 2. This would be a fairly small effect, however, and we do not see evidence of it in our data. As noted, there is

no significant difference in Configuration 4:0 between the amounts of bias (forward and backward) in Experiments 1 and 2. Also, the amount of overall overestimation in Experiment 2, compared with that in Experiment 1, was somewhat greater in Configuration 4:0 than in Configurations 4:4 and 4:2, whereas an increase in the perceived depth interval from the front of the platform (carpet) to the marker would be expected to decrease the amount of overestimation in the setting of the marker.

3. There may be some systematic errors in perceived distance related to the forward and backward displacement of the platform carrying the marker, but these errors would not explain the bias that we found. When the marker platform is raised above the ground (Configurations 2:2, 4:2, and 4:4), its surface is projectively compressed relative to the ground, so that distances along the platform may be underestimated relative to distances along the ground, as discussed earlier (Loomis & Philbeck, 1999; Meng & Sedgwick, 2001). When the platform is moved forward, for example, it occupies a greater proportion of the overall distance between the observer and the marker; thus the overall underestimation of the marker's distance might be expected to increase. To compensate for this underestimation, the marker would have to be moved farther away, relative to the condition in which the platforms are aligned. This hypothetical effect of compression on perceived distance is in the opposite direction to the bias effect that we observe: When the platform carrying the marker is moved forward, the setting of the marker also moves forward, although not as much. Furthermore, the hypothetical effect of

compression on perceived distance would be expected to be quite small on the basis of the amount of displacement of the platform, the difference in compression between the platform and the ground, and previous data (Loomis & Philbeck, 1999; Meng & Sedgwick, 2001). Finally, when the marker platform is *not* raised above the ground (Configurations 0:0, 2:0, and 4:0), moving the platform forward or backward has no effect on the overall compression and so would be expected to have no effect on perceived distance; yet we still find bias effects.

4. It is entirely arbitrary that we express the relative distance between the two platforms as the distance of the left platform minus the distance of the right platform (D_{L-R}). We could equally well express this relative distance as the distance of the right platform minus the distance of the left platform (D_{R-L}); if we did so, then Equation 1 would have to be changed to $P_R = P_L - D_{R-L}$ to compensate for the change in sign. The difference between these two forms (D_{L-R} and D_{R-L}) of Equation 1 is not substantive, but is only a difference in sign convention. Likewise, Equation 2 would become $P_R = P_L - k * D_{R-L}$, $k < 1$. In either the D_{L-R} form or the D_{R-L} form, Equation 2 makes exactly the same predictions for P_R , given a particular value of k .

(Manuscript received October 22, 2000;
revision accepted for publication February 5, 2001.)