Relative size perception at a distance is best at eye level

MARCO BERTAMINI, TYRONE L. YANG, and DENNIS R. PROFFITT
University of Virginia, Charlottesville, Virginia

Relative size judgments were collected for two objects at 30.5 m and 23.8 m from the observer in order to assess how performance depends on the relationship between the size of the objects and the eye level of the observer. In three experiments in an indoor hallway and in one experiment outdoors, accuracy was higher for objects in the neighborhood of eye level. We consider these results in the light of two hypotheses. One proposes that observers localize the horizon as a reference for judging relative size, and the other proposes that observers perceive the general neighborhood of the horizon and then employ a height-in-visual-field heuristic. The finding that relative size judgments are best around the horizon implies that information that is independent of distance perception is used in perceiving size.

Size constancy and size perception at a distance are old problems with theoretical implications for the nature of the perceptual process. Is perception of size derived inferentially from information about distance? Does perception of size rely on the general characteristics of the terrestrial environment in which we evolved? The current studies focus on the use of one source of information about the relative size of objects at a distance: the horizon line.

In this paper, we first consider the geometry of horizon information about size and discuss the various ways horizon information could be used for size perception. We then present a set of studies that examine how horizon information influences relative size perception. Our methodology avoids the complications that come with involving changes in perceived distance. We propose that, if the horizon is used as a reference for perceiving relative size, observers should better discriminate small height differences near this reference.

Before we can assess the role of horizon information about size, we need to consider the different ways in which an observer can use the horizon. For a standing observer viewing objects on the same ground plane, the explicit or implicit horizon line intersects objects at the observer's eye height, thereby specifying the absolute size of the object as a multiple of the observer's eye height (Sedgwick, 1973). This size information is independent of the distance of the object from the observer. Sedgwick (1973) was the first to observe that absolute size information is available in units of the standing observer's eye height, as is shown in the following formula:

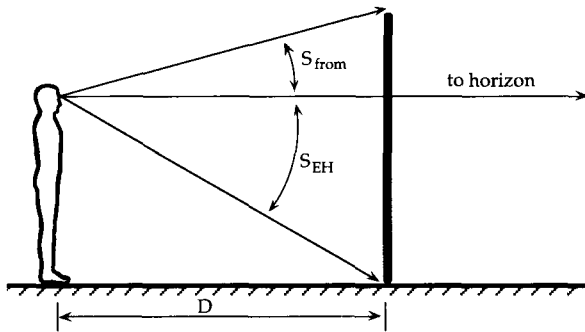$$H = \frac{\tan(S_{EH}) + \tan(S_{from})}{\tan(S_{EH})},$$

where H is object height as a fraction of eye height (sometimes called the horizon ratio), S_{EH} is the visual angle subtended by the object below the horizon—that is, the visual angle between the horizon line and the point of contact between the object and the ground—and S_{from} is the visual angle subtended by the object from the horizon (taken to be positive, if above, and negative, if below). Because the tangent function for small angles is close to linear, we can drop it and use, as an approximation for H, the following:

$$H \cong \frac{S_{EH} + S_{from}}{S_{EH}}.$$

Figure 1 provides a pictorial description of the terms in the formula. What makes the horizon special is the fact that it is effectively infinitely far away, and, therefore, the line connecting the point of observation with the horizon is parallel to the ground.[1] The distance between these two planes corresponds to the eye height, as is shown in Figure 1, and is a known constant for each individual. A reasonable approximation to infinity will give a reasonably good estimate of the horizon ratio H.

Absolute size can be obtained even if the observer changes his or her posture, as long as the observer can perceive and use information about the height of the eyes above the ground. The measure H would then specify height as a fraction of this new distance rather than of standing eye height. Recent data by Wraga and Neisser (1995) suggest that there are differences in the use of eye height information between postures. They concluded that eye height is used when people are standing or sitting but not when they are prone on the ground. Because of this

Figure 1. The horizon ratio is defined as $H = (S_{EH} + S_{from})/S_{EH}$. S_{EH} is the visual angle between the points where the horizon and the ground intersect the object. S_{from} is the visual angle between the horizon and the top of the object. S_{from} is zero when the object is at eye height, positive if the object is taller, negative if the object is smaller. For a given object and a given point of observation, the value of H is independent of the distance D.

relationship between posture and horizon line, in the rest of this article we use the term *eye level* to refer to the height of the observer's eye above the ground and the term *eye height* to refer specifically to standing eye level.

Extracting absolute size from horizon information requires three assumptions: (1) the *ground-as-support assumption*—the object lies on the ground; (2) the *common-ground assumption*—the observer and distant objects rest on the same flat ground; and (3) the *known-eye-level assumption*—the perceptual system registers and uses the distance between the eye and the ground to scale objects.

Although there is some evidence that people use eye height to judge absolute size (see, e.g., Mark, 1987; Warren & Whang, 1987), the relative role of eye height in conjunction with other sources of information (about the size of distant objects) and in different environmental conditions is currently not completely understood. For instance, Bingham (1993) has argued that observers rely more on knowledge about the shapes of common objects—such as trees—to scale the size of objects at a distance. He also argued that the conditions necessary to use the horizon ratio are difficult to satisfy: "the eye height hypothesis is controversial however, because eye height does not remain invariant as observers alter either their posture or the heights of the surfaces on which they rest. Eye height varies continuously as an observer adjusts from standing upright one minute, to kneeling or sitting the next, or as an observer locomotes along a trench, over a hill or up a set of stairs" (p. 1148).

Given the possible difficulty of satisfying these constraints, how might humans still use eye height information for perceiving size? We suggest that, when possible, the active perceiver may seek out situations conducive to the use of the horizon. We also suggest that the perceiver may be able to use horizon information to judge the relative sizes of distant objects. The situations in which the perceiver can use horizon information for relative size perception are more extensive than those for absolute size perception (Sedgwick, 1973, 1983).

Instead of absolute size, metric relative size can be gauged by the observer by comparing the horizon ratios from the previous equation for two objects. In this case, the known eye-level assumption (3) can be dropped (since the perceiver is no longer scaling objects to body size), and the common-ground assumption (2) can be relaxed (only the objects need to lie on a common horizontal surface, whereas the observer can be positioned anywhere). These relaxed assumptions would encompass situations such as viewing from within a trench or from up a set of stairs.

Furthermore, relative size information is preserved in pictures, in which the height above the ground of the center of projection (i.e., the height of the camera) is not explicitly known. Rogers and colleagues (Rogers, 1995; Rogers & Costall, 1983) have shown that people use the horizon to perceive relative size in pictures. Observers saw pictures of a pole and a horizon line. They tried to draw another pole of the same depicted size but farther away in the pictorial space (higher in the picture plane). Observers seemed to use the horizon ratio relation in their responses; they drew poles for which the ratio of the segment above the horizon to the entire pole was roughly the same as that for the target pole.

Rogers (1995) has pointed out that observers can use horizon information to perceive relative size in pictures, but there are reasons why this sort of relative size information also applies to nonpictorial situations. Even if a flat surface is a reasonable approximation for normal terrain, this is only true for limited regions of this space. Flat patches are intermixed with ramps, slopes, and higher or lower plateaus—for example, a person standing close to a creek will most likely be lower than the surrounding flat terrain. Topographical variations are more problematic for absolute size perception than for relative size perception. Furthermore, relative size information is not affected by postural change, which is essentially equivalent to raising or lowering the position of the terrain supporting the objects.

Finally, there is an even more basic use of horizon information for perceiving size. A binary decision between objects bigger than oneself and objects smaller than oneself can be based on the position of the top of an object with respect to the horizon line. In this case, the size of the object is not quantified, and there is only a discrimination between two categories.

What is common to all of these levels of obtaining size information from the horizon is that the observer needs to extract the location of the horizon line. That is, this location is necessary computationally, even if it has no phenomenal reality. Basing the computation of object size on a line other than the true horizon or misperceiving the location of the horizon line would result in errors in absolute, relative, and even categorical size judgments.

In a set of experiments that were similar to the original work by Sedgwick (1973), we used the same dependent variable: relative size judgments about two vertically oriented objects. More specifically, in our experiments, stationary observers judged which of two poles at a distance

was taller. The horizon line was never explicitly present. In four studies, observers viewed two poles placed at different distances in a hallway, with strong linear perspective implicitly specifying the location of the horizon. In another study (Experiment 4), observers viewed poles in a field. Observers judged which pole was physically taller.

If the implicit horizon line is used in judging relative size, size discrimination should peak near this reference and diminish above and below it. In other words, performance on a size discrimination task should vary not with the absolute height of the poles but with their absolute vertical distance from eye level. Moreover, since horizon information varies with observer height and posture, we predicted that, for the same objects, size discrimination would be better when the observer's eye level was near the height of the objects. For a person sitting on a chair, the horizon line is lowered, and size discrimination should be better for objects about as tall as sitting eye level. We tested this possibility in Experiment 1. Experiments 2 and 3 tested whether size discrimination performance would peak at eye level and progressively worsen as objects moved away from this reference. Experiment 4 examined the influence of the hallway ceiling on relative size perception. In Experiment 5, we tested the possibility that observers might be using an alternative height-in-visual-field (HIVF) strategy to make relative size judgments.

The essence of horizon information for perceiving size is the idea that the horizon provides size information that is independent of distance information. Consistent with this hypothesis, previous studies have shown that manipulating eye level or apparent eye level affects size perception (see, e.g., Warren & Whang, 1987; Wraga & Neisser, 1995). However, in theory, it is possible that manipulating eye level could influence perceived size not directly but through changes in perceived distance. In fact, there is some, at least phenomenological, evidence that changing eye level changes perceived distance (see, e.g., Carr, 1935, pp. 170–171). Our paradigm tests the use of horizon information for size perception without the complication that comes from possibly involving perceived distance.

EXPERIMENT 1

We investigated whether observers could use the implicit horizon as a reference line in judging the relative size of objects at different distances. Two conditions tested size discrimination in the neighborhood of two standard heights of the objects: short poles as high as the sitting eye level of the observer and tall poles as high as the standing eye level of the observer. Because the location of the horizon changes with posture, if the horizon is used as a reference, observers should better detect small size differences between the short poles when they are sitting and better detect small size differences between the tall poles when they are standing. To appreciate the relationship between the horizon and the height of the objects, refer to Figure 2. This line drawing traces a photograph of the hallway used in the experiment. In both

A   B
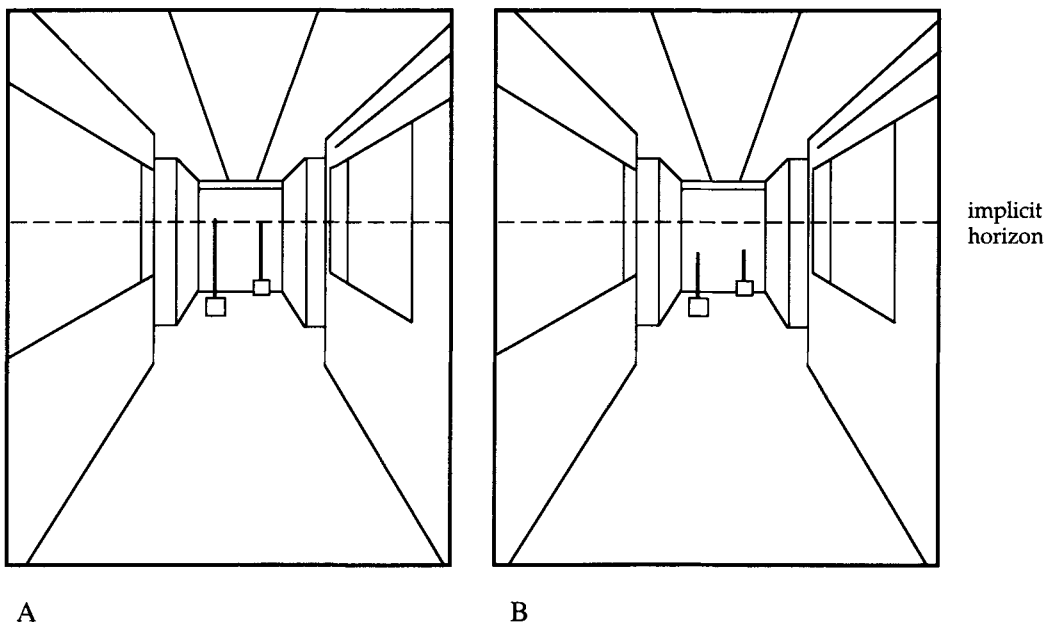
Figure 2. Two photographs were taken of the hallway as we used it in the experiment, and their lines were traced to show the geometry of the stimulus. The point of observation corresponds to the eye height of one of the authors (159.5 cm). In panel A, the right pole is 159.5 cm in height, and the left pole is 162.5 cm. In panel B, the right pole is 96 cm (60% of 159.5), and the left pole is 99 cm. Slight imperfections in the measures are due to the drawing process.

panels A and B, the point of observation corresponds to the standing eye level of one of the authors (159.5 cm). Panels A and B show, respectively, the tall poles condition and the short poles condition; in both cases, the front pole is 3 cm taller than the back one. In panel A, the top of the poles align with the implicit horizon, thus allowing an easy comparison of relative height (the left one is slightly above the horizon). In panel B, the tops of the short poles are distant from the horizon line, making for a more difficult comparison. Notice that the reverse would be true if the point of observation were lowered.

Method

Subjects. Ten University of Virginia students participated. They were naive with respect to the problem and the hypotheses until after the data were collected.

Design. A 2 (posture: standing vs. sitting) × 2 (pole height: tall vs. short) within-subjects design was used. The poles varied in height around either the observers' standing eye level (tall poles) or their sitting eye level (short poles).

Stimuli and Procedure. Each observer looked down a corridor and judged which of two poles was taller. The corridor was 1.83 m wide by 2.44 m high. The walls, tiles, and ceiling of the hallway provided plenty of linear perspective information. Pictures taken from one of the authors' eye level (159.5 m) were traced to create a line drawing and are shown in Figure 2.

Two white plastic PVC poles, 2.2 cm wide, stood upright at different distances from the observer. The far pole was 30.5 m away, and the near pole was 23.8 m away. They were separated from each other horizontally by 93 cm. Behind the poles, the end of the hallway was hidden by a black curtain hanging from the ceiling, 32 m from the observer. The size of this background was, therefore, 3.27° × 4.36° of visual angle. This was done for two reasons. The main reason was to avoid any reference line behind the objects; the other reason was to increase the visibility of the white poles. For the standing condition, the observer viewed the poles standing up. For the sitting condition, the observer viewed the poles sitting on a chair, with the head resting on a flat chinrest on a tripod. The tripod was set so that the observer's eye level was 60% of eye height.

For the tall poles condition, the standard pole was set at the observer's eye height. For the short poles condition, the standard pole was set at the observer's sitting eye level, at 60% eye height. Randomly, for half the trials, the left pole served as the standard, and, for half the trials, the right pole served as the standard. The other pole (the comparison pole) differed from the standard by -7, -5, -3, -1, $+1$, $+3$, $+5$, or $+7$ cm. In visual angles, an adjustment of 7 cm corresponds to 0.13° for the farther pole and 0.17° for the closer pole. In each of the four conditions (2 postures × 2 pole heights), observers made 16 judgments, for a total of 64 responses.

The procedure was the following. The observer stood at one end of the hallway, facing the wall opposite from the poles, while one of the experimenters adjusted the poles' heights and then hid in one of the doors of the hallway. At this point, the observer was allowed to turn around and look at the poles. After the judgment, the observer turned around again for the preparation of the next trial. The task of the observer was to judge which pole was taller and also to express his or her confidence in this judgment on a 10-point rating scale, in which 1 meant *not confident at all* and 10 meant *absolutely sure*. A second experimenter, standing by the observer, collected the answers on a data sheet.

Results and Discussion

For the tall poles, size discrimination was better when the observer was standing, and, for the short poles, size
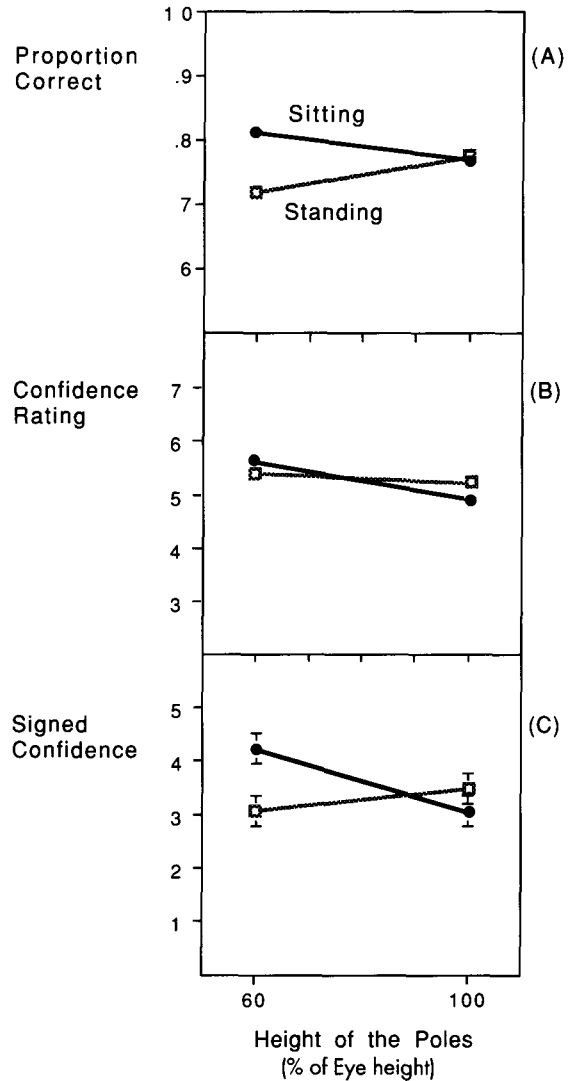
Figure 3. Data from Experiment 1. (A) Proportion correct for the sitting (closed circles) and the standing (open squares) conditions as a function of the standard height of the poles. (B) Average level of confidence on a 10-point scale. (C) Average of the positive (correct) and negative (incorrect) level of confidence. The maximum score on this derived variable is 10 and would be obtained only with correct judgments and maximum confidence on all trials. Bars are 95% within-subjects confidence intervals for the means based on the omnibus error term.

discrimination was better when the observer was sitting, suggesting a relationship between the location of the horizon and size discrimination performance. We computed proportions correct for the responses and average confidence. They are shown in Figures 3A and 3B, respectively. Then we computed a new derived variable by taking the confidence and coding it as positive (correct response) or negative (incorrect response). This new variable is bounded by the values 10 and -10. An average score of 10 means perfect performance: the observer was always correct and was always at the maximum level

of confidence. An average score of -10 means that the observer was always incorrect and was always at the maximum level of confidence. An average score of 0 means random performance: weighted correct and incorrect responses balance out. Thus, signed confidence provides information that is not available from separate analyses of percent correct and confidence alone, and the magnitude of the positive values is a measure of performance in the task. The means for the signed confidence variable are plotted in Figure 3C.

Performance as measured by percent correct or by signed confidence showed the same crossed interaction. The differences in overall level of confidence for the different conditions were very small, as is seen in Figure 3B, but had the same pattern as percent correct. Note that percent correct and signed confidence are measures of performance, whereas the level of confidence simply shows how people used the confidence rating on average.

Because of the analysis of variance (ANOVA) assumptions that require normal distributions, we did not analyze percent correct directly. Instead, we ran a repeated measures ANOVA on signed confidence. As was expected, the analysis found no main effect for posture (standing vs. sitting) and pole height (tall vs. short) but a significant interaction between the two $[F(1,9) = 7.107, p < .05]$. For the tall poles, performance was better when the observers were standing; for the short poles, performance was better when the observers were sitting. Figure 3C also shows 95% confidence intervals for the means based on the Loftus and Masson (1994) formula. Because this is a two-way repeated measures ANOVA, the standard error used to compute confidence intervals is the omnibus standard error based on the sphericity assumption.

EXPERIMENT 2

The outcome of Experiment 1 suggests that performance depends on the location of the horizon and that performance is higher when the object's height is close to eye level. However, the first experiment did not investigate how performance changes with object height above and below eye level. If the horizon is used as a reference for perceiving relative size, size discrimination should be easiest at the horizon and should get progressively worse farther from the horizon. In the second experiment, we explored the function of performance for a standing observer by using four sets of pole heights that were around 90%, 95%, 100%, and 105% of the observer's eye height.

Method

The procedure for Experiment 2 was identical to that for Experiment 1, except that, for all of the trials, the observer stood while viewing the poles. There were four pole conditions, in which the standard pole was set at 90%, 95%, 100%, and 105% of the observer's eye height. Sixteen University of Virginia students took part in this experiment. They were naive with respect to the problem and the hypotheses until after the data were collected.
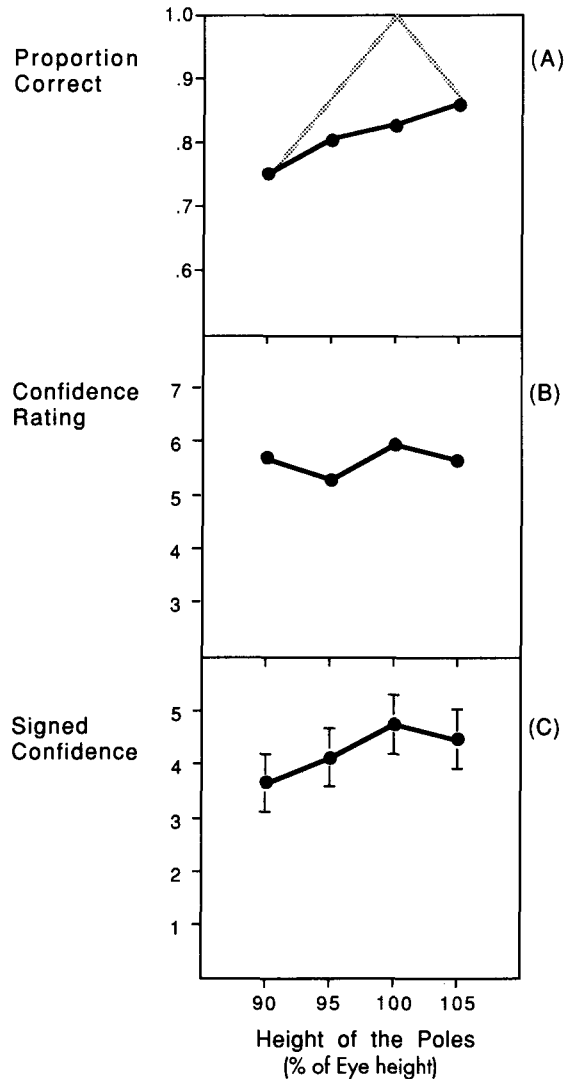
Figure 4. Data from Experiment 2. (A) Proportion correct as a function of the standard height of the poles. The gray line shows theoretical performance for an observer using a height-in-the-visual-field strategy. (B) Average level of confidence. (C) Average of the positive (correct) and negative (incorrect) level of confidence. Bars are 95% within-subjects confidence intervals for the means.

Results and Discussion

The analysis of the data was similar to that in Experiment 1. Mean performance on proportion correct (panel A), average confidence (panel B), and signed confidence (panel C) is plotted in Figure 4. This figure shows that performance on size discrimination increased with object height, if measured only by proportion correct. The more sensitive measure based on the signed confidence, however, shows that performance was better for objects at eye height than for objects at closely neighboring heights. The differences in overall level of confidence for the different conditions were small and have no clear pattern, as is seen in Figure 4B.

We ran a repeated measures one-way ANOVA on signed confidence and found a significant effect of condition $[F(3,42) = 3.04, p < .05]$. Post hoc least significant difference (LSD) tests found differences only between the 90% and the 100% and 105% conditions. Therefore, the tests cannot conclusively say that statistically the peak of the function is at 100%. More evidence about the shape of the function will come from the following experiments.

As mentioned before, there is more than one way to use the horizon line in size judgments. One possibility is that this line simply allows a categorical discrimination between what is above and what is below the horizon. More specifically, for objects about as high as eye level, it is possible to simply perceive as taller the one whose top is higher in the visual field. To describe what kind of performance is predicted by this strategy, we computed for each observer's height the level of percent correct based on always choosing as higher the pole whose top was higher in the visual field. This theoretical performance is plotted in Figure 4 as a gray line and should be compared to the observed performance. Around the horizon, this strategy leads to perfect performance and falls symmetrically for standard poles below and above the horizon. In this experiment, the two lines are similar, with the observed performance being below the theoretical, as should be expected from subject error. The HIVF strategy could therefore explain the observers' responses in this experiment. However, during debriefing and questioning after the experiment, none of the observers reported looking at the location of the top of the pole. If this strategy was adopted, it was not done consciously.

EXPERIMENT 3

In Experiment 2, we used poles, the height of which ranged at most by 10% from the observer's eye height, and found that small size discriminations were better in the neighborhood of eye height. However, the function peaks at 100%, or it might only level off for higher values. The third experiment explores a larger range of heights with the same setting and methodology. We expected performance to drop both above and below the value of 100%.

Method

The procedure of Experiment 3 was similar to the procedure of Experiment 2. However, the standard pole was set at 80%, 100%, 110%, and 120% of the observer's eye height. Sixteen University of Virginia students took part in this experiment. They were naive with respect to the problem and the hypotheses until after the data were collected.

Results and Discussion

The analysis was similar to that in Experiments 1 and 2, and mean performance on proportion correct, confidence, and signed confidence is plotted in Figure 5. We ran a repeated measures ANOVA on the signed confidence and found no significant main effects or interactions.

Looking at the means in Figure 5, it appears that size discrimination performance peaked at 110%. Interpreta-
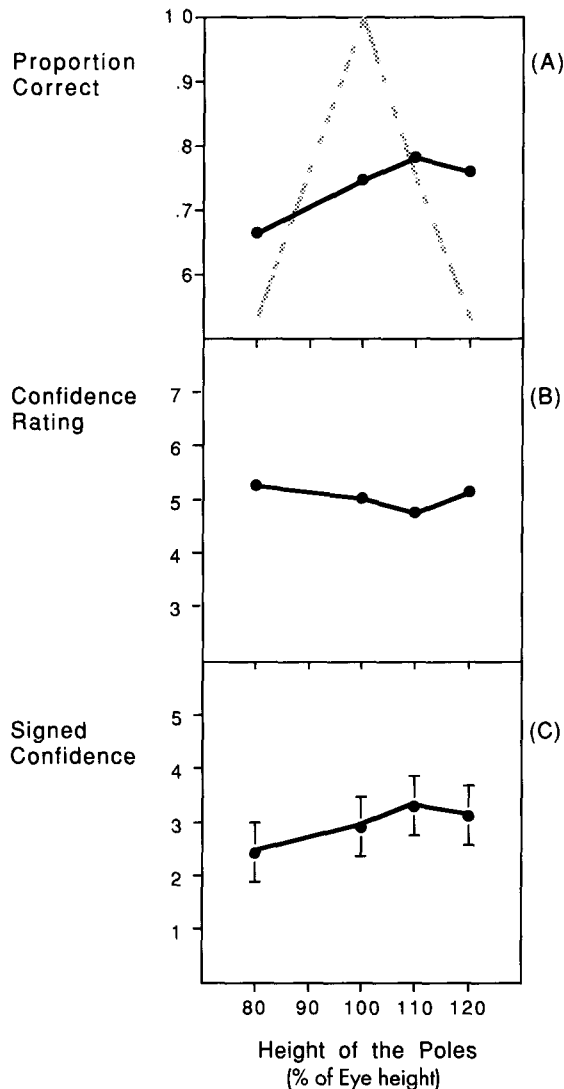

Figure 5. Data from Experiment 3. (A) Proportion correct as a function of the standard height of the poles. The gray line shows theoretical performance for an observer using a height-in-the-visual-field strategy. (B) Average level of confidence. (C) Average of the positive (correct) and negative (incorrect) level of confidence. Bars are 95% within-subjects confidence intervals for the means.

tion is difficult because of the high variability, as can be seen from the 95% confidence intervals. However, the peak of the function at 110% is not in agreement with our prediction, even though the function shows a drop at both ends, as predicted. Tentatively, we can say that the function seems nonmonotonic, but its shape requires more data and may also depend on the specific landmarks available in the environment. A possible explanation for the performance peak at 110% is that the poles were higher than those in Experiment 2. With tall poles, the ceiling of the hallway could act as a reference for relative size judgments in a manner similar to that of the horizon. In a situation in which horizon information is more difficult to extract—for instance, when only implicit horizon infor-

mation is available—the visual system may rely on other sources of information.

The theoretical performance based on an HIVF strategy is plotted in panel A of Figure 5. The gray line shows what is predicted for observers who systematically pick as higher the pole whose top is higher in the visual field. Because lower and higher (80% and 120%) standard poles were used, this strategy now predicts performance very close to chance for these conditions. However, the observers show an average performance in these conditions around 70% correct. It is highly unlikely that the observers in this experiment adopted an HIVF strategy.

EXPERIMENT 4

Experiments 2 and 3 explored performance on size discrimination as a function of object height. We found that the function is highest around the value of eye height (100%) but could not pinpoint the peak. In contrast with the results of Experiment 2, the peak in Experiment 3 was at the 110% eye height condition. We hypothesize that this difference was a consequence of an interaction between two factors: although the horizon line improves performance around eye level, size discrimination may also improve around other environment-specific references, such as the hallway ceiling. Experiment 4 was designed to replicate the shape of the function found in Experiment 3 and to test our prediction that, in an outdoor setting where there is no ceiling, performance should peak at eye level.

Method

The procedure was similar to that in Experiment 3: The standard pole was set at 80%, 100%, 110%, and 120% of the observer's eye height. Half of the subjects were run indoors, and half were run in a grassy field outdoors, which had bushes, trees, and vines in the background. The ground was flat, and the field was used occasionally for soccer. The only difference between the indoor setting and the outdoor setting in terms of procedure was that the experimenter could no longer hide in a doorway. Instead, he moved several meters away before the observer looked at the poles. Thirty-two University of Virginia students took part in this experiment. They were naive with respect to the problem and the hypotheses until after the data were collected.

Results and Discussion

The analysis was similar to that of the previous experiments, and mean performance on proportion correct, confidence, and signed confidence are plotted respectively in panels A, B, and C of Figure 6. Because the comparison between overall level of performance in the indoor and outdoor conditions was not critical, the 95% confidence intervals of panel C are based on two separate one-way repeated measures ANOVAs.

A two-way mixed ANOVA on signed confidence revealed an effect of condition [$F(3,90) = 2.797, p < .05$], a nonsignificant effect of setting [$F(1,30) = 3.943, p < .056$], and a significant interaction between the two [$F(3,90) = 3.219, p < .05$]. LSD post hoc tests for the interaction found that, for the indoor setting, the 80% con-
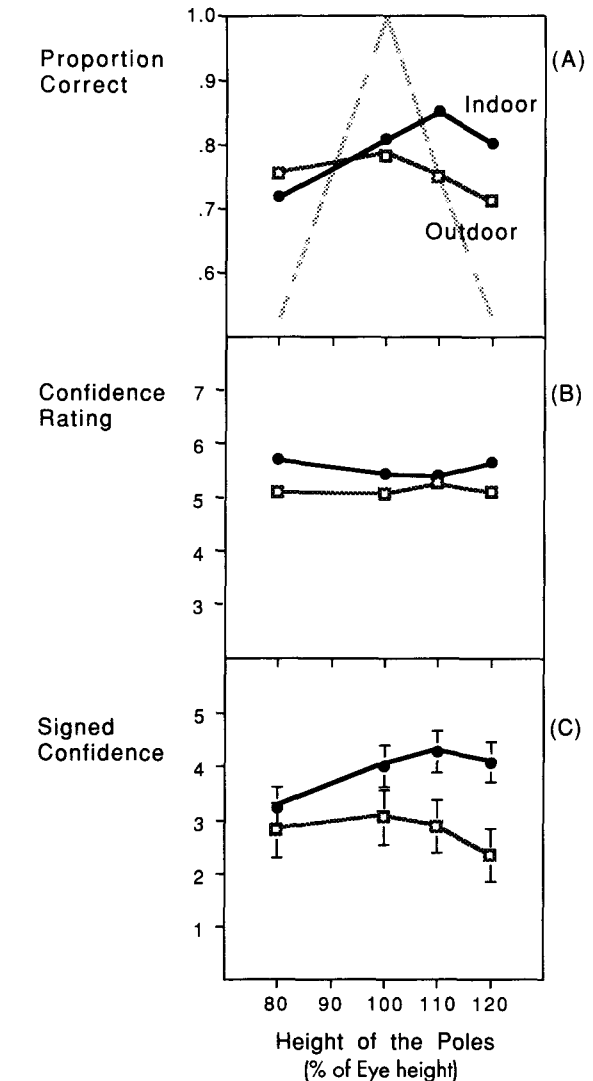
Figure 6. Data from Experiment 4. (A) Proportion correct for the indoor (closed circles) and the outdoor (open squares) conditions as a function of the standard height of the poles. The gray line shows theoretical performance for an observer using a height-in-the-visual-field strategy. (B) Average level of confidence for the indoor and outdoor conditions. (C) Average of the positive (correct) and negative (incorrect) level of confidence. Bars are 95% within-subjects confidence intervals for the means based on two separate within-subjects ANOVAs.

dition was significantly smaller than the 100%, 110%, and 120% conditions. For the outdoor setting, the 100% was significantly greater than the 120% condition.

The means in Figure 6 suggest that size discrimination was better for objects around eye level (100% condition) for the outdoors condition, but, for the indoors condition, performance peaked at 110%. This is in agreement with our expectation and with the results of Experiment 3. When high poles are close to the ceiling, there is a new reference that can be used to make size judgments; this would explain the difference between indoors and outdoors (there is no ceiling in the outdoor setting) and also

the fact that performance peaked at 100% in Experiment 2 (shorter poles). We conclude that people performed the task in a similar manner under both settings, by basing their judgments on the implicit horizon line.

The theoretical performance based on an HIVF strategy is plotted in panel A of Figure 6. This function is the same shape as that in Figure 5; again, for the 80% and 120% standard pole conditions, the observed performance is much higher than what is predicted by an HIVF strategy.

EXPERIMENT 5

We have found that size discrimination performance peaks for objects around eye level. We have proposed that observers' use of the horizon as a reference to make relative size judgments could produce this peak. We also considered an alternative perceptual strategy in which observers select as physically taller the object whose top is higher in the visual field, which would result in a similar pattern of performance. An observer using this HIVF strategy would correctly judge relative physical size 100% of the time for objects around eye level. Geometry predicts that, using the HIVF strategy, performance would also progressively worsen as objects became larger or smaller than eye height. The HIVF hypothesis is attractive, because it circumvents possible limitations of the visual system in precisely locating the horizon and making the fine size discriminations that are required in our experiments. Previous experiments on perception of eye level have shown that observers can localize the horizon but that there is a systematic error below the true horizon of about 0.13° (MacDougall, 1903) or 0.29° (Stoper & Cohen, 1986). In the latter study, the standard deviation was 1.03°, and the unsigned deviation from the mean was 0.80°. The systematic error is larger in the dark, and, presumably, there is a continuum from a completely dark room (largest error) to a lit room, to a long hallway, up to the condition where the horizon is present explicitly (no error).

The pattern of performance in Experiments 3 and 4 in some conditions exceeded the predicted level of performance with the HIVF strategy. This is consistent with the hypothesis that observers did not solely choose the pole that was higher in their visual field. The goal of Experiment 5 was to see what is the actual level of performance of observers asked to use the HIVF strategy. To estimate how well performance could be accounted for by the HIVF strategy, a new set of observers in Experiment 5 tried to pick which pole was higher in their visual field, and their judgments were scored as correct if the pole they picked was the physically taller pole.

Method

The procedure was similar to the procedure in Experiment 3. The standard pole was set at 80%, 100%, 110%, and 120% of the observer's eye height. However, instead of making objective relative size judgments, the observers indicated which pole was higher in their visual field. Eight University of Virginia students participated in this experiment. They were naive with respect to the problem and the hypotheses until after the data were collected.

Results and Discussion

We first computed how well observers chose the pole that was higher in the visual field and then computed how well observers would perform if their judgments were scored for choosing the pole that was physically taller. Panel A of Figure 7 shows the overall level of performance in picking the higher top in the visual field. It is clear that the task is not trivial (percent correct was between 70% and 95%) and becomes harder for higher poles. This decline could be explained by a preference for choosing the pole farther away, because this tends to be the correct
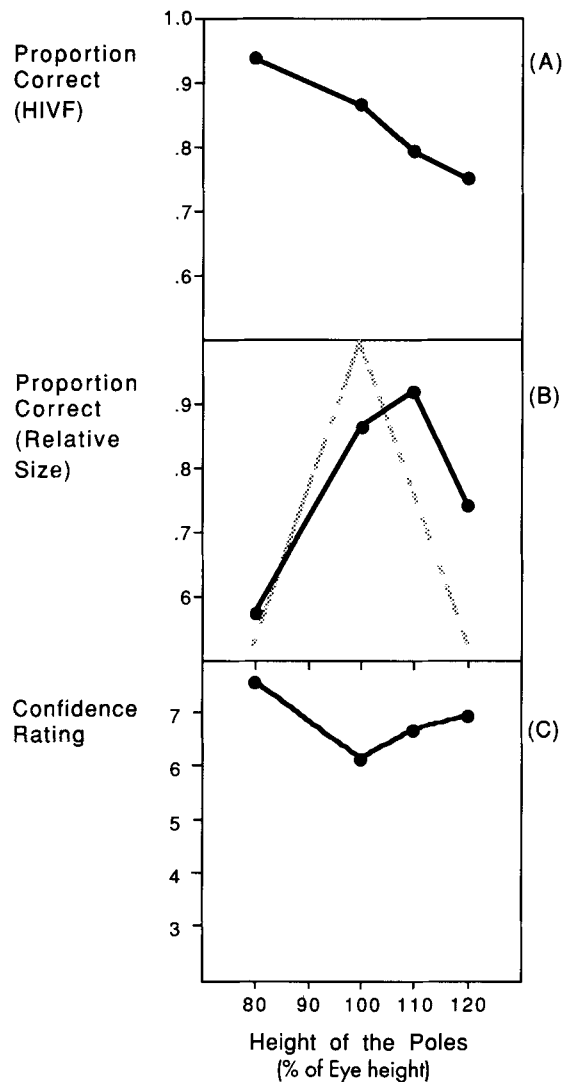
Figure 7. Data from Experiment 5. (A) Proportion correct as a function of the standard height of the poles. (B) Proportion correct computed on the basis of how often the observers chose the physically taller pole (which they were not asked to do). The gray line shows theoretical performance for an observer using a height-in-visual-field strategy. (C) Average level of confidence.

response for short poles and the wrong response for taller poles.

Panel B of Figure 7 shows how well observers would perform if their judgments were scored for choosing the pole that was physically taller. That is, on the basis of which pole was judged by the observers to have a higher top in the visual field, we computed how often that choice corresponded to the physically taller pole. This performance should be compared to the theoretical performance, shown by the gray line in panel B. The shape of the two functions are similar but differ largely in the 110% and 120% conditions. Surprisingly, in these conditions, the proportion correct in the size task is higher than what is predicted on the basis of perfect performance with the HIVF strategy. This suggests that, when the observers made an error, it was from choosing the pole that was physically taller more often than from choosing the one that was shorter.[2]

The goal of this experiment was to compare the performance of observers who were instructed to look at the HIVF with the performance of observers in previous experiments (3 and 4), where they were instructed to judge the physical size. To make this comparison, the signed confidence data were entered in a new ANOVA. The factors were condition (80%, 100%, 110%, and 120%) and task (size in the combined previous experiments vs. HIVF in Experiment 5). There was a significant effect of condition [$F(3,114) = 15.62, p < .0001$] and no significant difference between tasks [$F(1,38) = 4.00$]. But there was a significant interaction between condition and task [$F(3,114) = 5.73, p = .005$]. Because signed confidence in the HIVF task can only be computed on the basis of the assumption that confidence in judging the pole that is higher in the visual field is the same as judging which pole is physically taller, we also ran the same ANOVA with the percent correct data. This analysis yields the same pattern of effects.

According to an LSD post hoc test, in the 100% and 110% conditions, performance was significantly better in the HIVF task than in the combined Experiments 3 and 4. The LSD post hoc test from the percent correct data confirms that, in the 100% condition, performance was significantly better in the HIVF task (91% correct) than in combined Experiments 3 and 4 (79% correct). Performance in the 80% eye height conditions, on the other hand, was significantly worse in the HIVF task (63% correct) than was the combined performance from Experiments 3 and 4 (71% correct). Overall, this confirms the pattern seen in Figure 7B, where performance was very low in the 80% condition and surprisingly high in the other conditions. The explicit use of the HIVF strategy alone does not model performance in previous experiments very well.

When asked at the end of the trials, the subjects in the previous experiments (2, 3, and 4) did not report only attending to the relative tops of the poles (heights in visual field) when judging relative size. Also, performance using the HIVF strategy differed from the objective size discrimination performance in Experiments 3 and 4. Nevertheless, the current data do not allow us to completely discount the possibility that observers may use the HIVF strategy at least some of the time. Although this HIVF strategy could work well around eye height, performance would be poor if observers were to adopt this strategy to compare very short poles (for which the farther pole would always be higher in the visual field) or to compare very tall poles (for which the closer pole would be higher in the visual field). An effective use of this HIVF strategy would require that observers employ this strategy selectively, using it only when objects are about as tall as the horizon.

GENERAL DISCUSSION

We have argued that observers use the horizon in perceiving the relative size of distant objects. The current paradigm tests this possibility without the complications that come with possibly involving perceived distance. The finding that relative size perception is best around eye level reflects this use of the horizon. This peak in size discrimination performance moves with postural change. Experiment 1 showed that, for a standing observer, size discrimination is better at standing eye level; for the sitting observer, size discrimination performance is better at sitting eye level. Using objects of a greater range of sizes, Experiments 2, 3, and 4 provide evidence, especially if considered together, that size discrimination performance is a nonmonotonic function of the height of the objects. All plots of mean performance (see Figures 4, 5, and 6) peak near eye level, but we were not able to pinpoint the exact location of this peak because of variability. This general pattern of performance holds both for architectured settings with plenty of linear perspective information and for outdoor settings.

This peak in performance is consistent with our hypothesis that observers use the implicit horizon as a reference for perceiving the relative size of objects. Size discrimination becomes more difficult, the farther objects are from this reference. However, other processes may also account for this pattern of performance. A strategy of selecting as physically taller the object that is higher in the visual field would also yield a peak in performance at eye level.

Experiment 5 examined the actual performance of observers instructed to use this HIVF strategy. The observers who used this strategy also performed best around eye level, but they also performed better in the 100% and 110% eye height conditions and significantly worse in the 80% condition than the observers from our other experiments who had made judgments about relative objective size. Given that the observers making objective size judgments performed differently than the observers using the HIVF strategy, we cannot, at this time, rule out the possibility that observers use the HIVF strategy at least some of the time. However, for objects far away from the horizon, performance with the HIVF strategy would be at chance. Thus, an effective use of the strategy would require at least roughly localizing the general area of the horizon before implementing this heuristic. If observers changed their strategies depending on the conditions and,

for instance, adopted a HIVF strategy for the 100% standard pole condition, this was not supported by the verbal reports after the experiment. None of our subjects reported adopting a different strategy for this condition.

In short, the peak in size discrimination performance around eye level may be a consequence either of comparing horizon ratios or of comparing objects' HIVFs. A strategy of comparing horizon ratios portrays the human perceiver as a geometer. This strategy implies that the visual system perceives, to some level of accuracy, the location of the horizon. This information could come from an explicit, visible horizon or be implicitly indicated by linear perspective, by a discontinuity in optic flow, or from the limit in optical compression of ground texture. An implication of this hypothesis is that, in a more naturalistic situation, the active perceiver may try to align eye level with one of the objects (e.g., by tiptoeing or crouching) to better judge their relative sizes. Some informal observations in pilot experiments suggest that some observers do, in fact, do this. However, there are clear limits to the precision of size perception that is based on the horizon ratio: There are systematic errors in locating the implicit horizon, there are strong assumptions not always met in the natural environment, and there are errors that would ensue from measurement errors. Finally, the evidence from Experiments 3 and 4 also suggests that other reference objects in the environment—such as the ceiling—can play a role in perceiving size.

The second strategy portrays the perceiver as more of a heuristical beast. For objects that are roughly body-sized, observers use a HIVF strategy, which would be most effective for objects very close to eye height. Observers overgeneralize this strategy by applying it to objects that are only roughly around eye height. The HIVF strategy is one that also works best for objects near eye level and less well as objects deviate in size from eye level. Effective use of this strategy only requires that observers perceive the rough neighborhood of the horizon. With this strategy, observers do not compute horizon ratios at all but only attend to the relative HIVFs of the tops of objects around eye level. If observers in our experiments did use this strategy, one might wonder whether observers would use the horizon ratio if an explicit horizon were available, given that observers use the horizon ratio to judge relative size in pictures (Rogers, 1996).

The current experiments show that relative size perception is best at eye level. The relationship between the horizon and the objects to be perceived influences the perception of relative size. This relationship implies the use of information that is independent of perceived distance in the perception of the relative size of objects.

REFERENCES

BINGHAM, G. P. (1988). Task-specific devices and the perceptual bottleneck. *Human Movement Science, 7,* 225-264.

BINGHAM, G. P. (1993). Perceiving the size of trees: Form as information about scale. *Journal of Experimental Psychology: Human Perception & Performance, 19,* 1139-1161.

CARR, H. A. (1935). *An introduction to space perception.* New York: Longmans, Green.

LOFTUS, G. R., & MASSON, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review, 1,* 476-490.

MACDOUGALL, R. (1903). The subjective horizon. *Psychological Review Monograph Supplements, 4* (1, Whole No. 17), 145-166.

MARK, L. S. (1987). Eyeheight-scaled information about affordances: A study of sitting and stair climbing. *Journal of Experimental Psychology: Human Perception & Performance, 13,* 361-370.

ROGERS, S. (1995). Perceiving pictorial space. In W. Epstein & S. Rogers (Eds.), *Perception of space and motion* (2nd ed., pp. 119-163). San Diego: Academic Press.

ROGERS, S. (1996). The horizon-ratio relation as information for relative size in pictures. *Perception & Psychophysics, 58,* 142-152.

ROGERS, S., & COSTALL, A. (1983). On the horizon: Picture perception and Gibson's concept of information. *Leonardo, 16,* 180-182.

SEDGWICK, A. (1973). *The visible horizon: A potential source of visual information for the perception of size and distance.* Unpublished doctoral dissertation, Cornell University.

SEDGWICK, A. (1983). Environment-centered representation of spatial layout: Available visual information from texture and perspective. In J. Beck, B. Hope, & A. Rosenfeld (Eds.), *Human and machine vision* (pp. 425-458). New York: Academic Press.

SEDGWICK, A., NICHOLLS, A., & BREHAUT, J. (1995, May). *Perceptual interaction of surface and depth in optically minified pictures.* Poster presented at the annual meeting of the Association for Research in Vision and Ophthalmology, Fort Lauderdale, FL.

STOPER, A. E., & COHEN, M. M. (1986). Judgments of eye level in light and in darkness. *Perception & Psychophysics, 40,* 311-316.

WARREN, W. H., & WHANG, S. (1987). Visual guidance of walking through apertures: Body-scaled information for affordances. *Journal of Experimental Psychology: Human Perception & Performance, 13,* 371-383.

WRAGA, M. J., & NEISSER, U. (1995, July). *The effect of viewer's posture on perceived object size.* Poster presented at the 8th International Conference on Perception and Action, Marseille.

NOTES

1. The true horizon on our round planet is technically not infinitely far away, and its distance depends on eye level. For instance, for a person about 2 m tall, the horizon is only about 5 km away. However, the distance of the horizon is still very large, as compared to the height of the average person (Michelangelo Fluckiger, personal communication, April 1996).

2. In this experiment, the observers judged a proximal property of the stimulus. This assumes access to geometrical information of the two-dimensional projection in the visual field. As has been observed before (Sedgwick, Nicholls, & Brehaut, 1995), in such cases people may in fact be affected by the layout and the three-dimensional information, which they are asked to ignore. It is, therefore, possible that our observers found it difficult to ignore which pole looked taller in judging which top was higher in the visual field. This experiment, however, was not designed to test this possibility, and it is possible that the design itself created this higher performance as an artifact. Suppose that people tried to balance left and right responses (i.e., farther and closer poles); in the case of very short or very tall standard poles, this may lead them to choose the physically taller pole, because, in those trials, the tops are closer together in the visual field (for geometrical reasons). For example, for short poles, the correct answer would have been almost always "farther pole." If subjects tried to say "closer pole" sometimes, it is likely that they did so for conditions in which the closer pole was physically taller (because this brings the tops closer together in the visual field, and, therefore, it is less obvious that the farther one is higher in the visual field). This artifact should have increased the performance on the very tall and very short poles conditions of Figure 7B. Therefore, it may explain the high performance in the 100% and 120% conditions. It is at odds, instead, with the very low performance in the 80% condition.