Downhill slopes look shallower from the edge

Zhi Li
Department of Psychology, Swarthmore College, Swarthmore, PA, USA

Frank H. Durgin
Department of Psychology, Swarthmore College, Swarthmore, PA, USA

A dramatic failure of orientation constancy is documented in the perception of downhill slopes. Contrary to naïve expectation, steep downhill slopes look shallower from the edge than they do from back from the edge. Three experiments document and quantify this failure of constancy for real and virtual surfaces using a variety of dependent measures. Two additional studies document overestimation of both non-visually perceived head pitch and perceived gaze declination. A model of orientation constancy failure is fit to the data that combine exaggerations in perceived gaze declination with exaggerated scaling of perceived optical slant. These findings support a functional scale-expansion model of error in slope perception.

Keywords: slope perception, geographical slant, orientation constancy, frontal tendency, perceived gaze direction, perceived head orientation, virtual reality

Citation: Li, Z., & Durgin, F. H. (2009). Downhill slopes look shallower from the edge. *Journal of Vision, 9*(11):6, 1–15, http://journalofvision.org/9/11/6/, doi:10.1167/9.11.6.

Introduction

Hills generally look much steeper than they are (Kammann, 1967; Proffitt, Bhalla, Gossweiller, & Midgett, 1995), but there is disagreement about whether the cause of this error reflects a lack of information, a simple bias, or the behavioral potential of embodied perceivers (Bridgeman & Hoover, 2008; Gibson & Cornsweet, 1952; Ooi, Wu, & He, 2006; Proffitt et al., 1995). Here we report a new phenomenon concerning the perception of slopes viewed from the top that may help to adjudicate among some of the theories presently in the literature. The phenomenon is simple, but striking: The apparent downward slope of a hill, or even a small ramp or sloped surface, appears steeper when one stands back from the edge of the hill or surface.

We first observed this phenomenon while walking on a path along a wooded slope on our campus. How, we wondered, could a (down-sloping) hill of some 18 degrees appear to be nearly 45 degrees, when inclining the head by such an amount would clearly falsify this belief? In fact, as we stood back from the edge of the path to attain a parallel view along the hill surface, we observed that the hill appeared steeper, not shallower.

Bridgeman and Hoover (2008) have argued that the slope of hills viewed from the bottom depends on the distance to the portion of the slope that was judged. Whereas Proffitt et al. (1995) reported that hills appear steeper when viewed from the top than from the bottom, our observation that perceived slant from the top varies dramatically with viewing position suggests that the categorical comparison of “top” and “bottom” viewing is unwarranted. Unlike Bridgeman and Hoover, our account emphasizes the role of direction of gaze rather than viewing distance, though both factors likely affect the recovery of geographical slant.

In theory, perceived surface orientation might be regarded as resulting from a geometric combination of an estimate of gaze orientation (whether based on non-visual proprioception or on what Gibson, 1966, referred to as visual proprioception) and an estimate of surface orientation with respect to the line of gaze, known as optical slant (Gibson & Cornsweet, 1952; Sedgwick, 1986). Changing one’s position with respect to the top of a hill changes both the orientation of gaze required to view the hill surface and the resulting optical slant angle between the surface and the line of gaze. We will suggest that systematic biases in the perception of both of these variables may be sufficient to account for the fairly dramatic failure of surface orientation constancy we have observed.

In three of the experiments below, we document this failure of constancy for down hill surface orientation using both real and virtual surfaces. Two other experiments measure biases in the non-visual proprioception of head orientation and the perceived direction of gaze. We show that our slant perception data are well described by a simple theory relating geographical slant misperception to the misperception of gaze orientation and then describe in more detail how these findings bear on current theories of surface orientation perception.

doi: 10.1167/9.11.6 Received March 14, 2009; published October 6, 2009 ISSN 1534-7362 © ARVO

Experiment 1: Visual matches of small wooden slopes

Although the phenomenon was first observed on a large outdoor surface, we have observed that small surfaces also appear to dramatically change their orientations depending on one's point of regard. To document this basic failure of surface-orientation constancy without depending on verbal estimates of slope, we had naïve participants adjust a small board at their feet until it appeared to match the orientation (i.e., steepness) of a reference ramp just in front of them or a short distance away. If the matches produced when the reference ramp was close were shallower than those produced when it was far, we could conclude that there had been a failure of surface constancy consistent with our outdoor observations.

Methods

All participants in the experiments reported here had normal or corrected-to-normal vision, received instruction rather than training in the tasks, and were uninformed of the hypotheses. The participants in this experiment were 55 undergraduates who participated for course credit or pay. Each participant made two adjustments in one of six conditions representing the combination of three different reference angles with two different viewing distances. There were 8 to 10 participants in each cell.

The reference slope was a rectangular piece of unpainted plywood 1.02 m wide and 0.76 m long that was propped up at one of three angles (16.5, 24.3, or 32.1 degrees) on a low-contrast linoleum floor. The matching slope was a smaller board (0.5 m long and roughly 0.5 m wide) whose sides were irregularly curved to discourage attempts to use linear perspective (see Figure 1). This board was hinged from below to the floor at the far side; the front edge could be lifted using a handle held at waist level that was attached by fishing line. The produced slope on each trial was measured with an inclinometer. Each participant produced two consecutive estimates in the same condition, returning the matching board to the floor between estimates. For the steepest incline, participants and matching board stood on a platform that elevated them by 7.5 cm so that they would be the same distance above the near edge of the reference slope as the participants in the middle slope condition. This ensured that shorter participants would be able to see the surface of the steep slope in the far condition clearly.

Results

Mean matched angles are shown in Figure 2. As predicted, a 2 (distance) $\times$ 3 (slope) ANOVA showed that the matching slope was set higher when the reference slope was farther away, $F(1, 49) = 5.88$, $p = .019$. Post hoc comparisons with Bonferroni correction showed that the effect was reliable for the two steeper slopes, but not for the 16.5-degree slope.

Experiment 2A: Verbal estimates of virtual slopes

Although Experiment 1 successfully documented the basic failure of downhill slope constancy, we sought to parametrically vary surface orientation in a visual context more similar to outdoor viewing, where surfaces are larger. For this purpose, we conducted an experiment measuring the perception of surface orientation in an immersive virtual environment. We used a virtual reality system that rendered rich three-dimensional ground textures with rapid updating, low lag, and accurate projective geometry.

Figure 1. Depiction of (between-subject) near and far positions for perceptual matching task of Experiment 1. The surface of the adjustable board was planar but had irregularly curved sides.

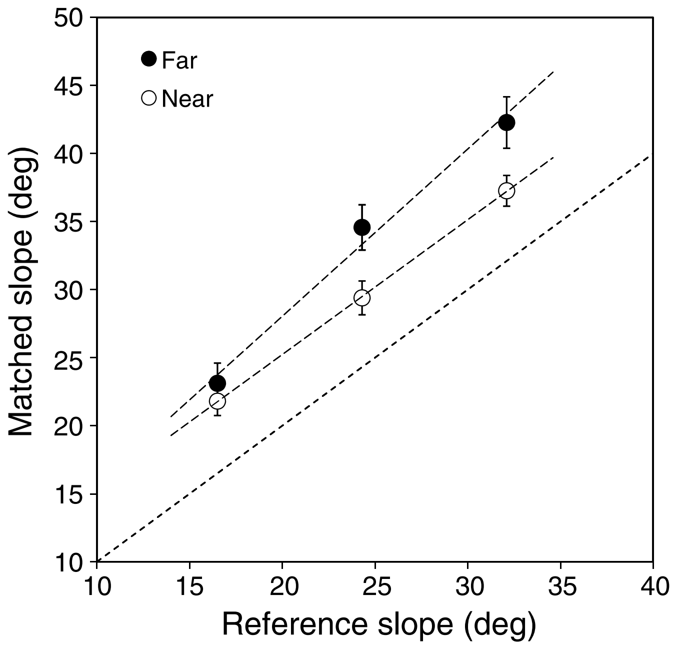

Figure 2. Results of Experiment 1. Matches of an adjustable slope to a fixed slope at one of two distances. Standard error bars are shown.

Methods

Participants were 24 students who had not been in previous experiments on slope perception.

Stereo graphics were rendered at 60 Hz to an nVisor (1280 × 1024) head-mounted display (HMD) using Virtools 4.1 and an nVidia 9800 GTX graphics card. Low-lag 6 DOF optical head tracking was provided by a HiBall system. A software shader compensated for the pincushion distortion of the optics in real time, providing a calibrated rectangular view of 42 × 32 degrees of visual angle. (Uncorrected optics would have produced motion artifacts during head rotations that could make scenes appear non-rigid.)

The virtual environment depicted a scene in which participants stood in a grassy walled-in area with Japanese maple trees to the left and right. (This scene was used to permit participants to become accustomed to the head-mounted display.) A wide aperture in the wall in front of participants was bordered by marbleized spheres on each side and a large cylinder at the base to minimized visually available orientation information. The sides of the aperture were usually not visible when looking straight down the hill.

The scene through the aperture depicted a lake 14.5 m below the level of the grass on which the observer stood. The lake extended 90 m to a distant land that defined the horizon. On each trial, a slope depicted as covered with irregularly spaced stones (each about 5 cm across) was visible through the aperture. The slope was 7.8 m long and extended to the left and right as far as could be seen through the aperture. A white ball 7.5 cm in diameter stood (on an unseen golf-tee) on the slope and was used as a reference mark during judgments. The horizontal grass surface was not normally visible when looking at the ball. Two views of the virtual environment are shown in Figure 3.

The task was to provide a verbal estimate of the apparent geographical slant of the hill in the vicinity of the ball. Most educated adults have a good understanding of verbal units of slope. To discourage categorical responding and reliance on memory, we informed our participants (truthfully) that they would never see the same simulated orientation twice and we encouraged them to give estimates that were more precise than the 5-degree increments that people naturally prefer. These instructions encourage careful numeric estimation.

Each of the participants made 30 judgments of slope across three conditions. In the main conditions, the ball was located 2 m down the sloped surface but the viewer was simulated as being either at the edge of the slope (Near) or 1 m back from the edge (Far). As a control for viewing distance, an additional condition was included in which the ball was located only 1 m down the sloped surface while the observer was simulated as being back 1 m from the edge (Far-control). Condition was blocked (10 trials in each of three conditions) with six possible condition orders distributed equally across the 24 participants. We used slopes that ranged in 1-degree increments from 5 to 34 degrees, such that every third slope was presented in one of our three viewing conditions. Participants always did the middle set of angles (6–33 degrees) in the first block and then the lower angles (5–32) and finally the higher angles (7–34) in the final block. We later reduced the data by local linear interpolation and eliminated the outermost slopes at each end of the range because pilot testing had shown that judgments of the extremes of any presented range were more variable. By this process, we arrived at individual estimates for eight slopes (9, 12, 15, 18, 21, 24, 27, and 30 degrees) in each of the three conditions for each observer. We also recorded head orientation information for each trial.

Before each trial, observers were required to hold their head upright while the display was blank and close their eyes; a baseline head orientation was recorded and the display turned on. Observers thus began each trial looking forward at the horizon and then moved their gaze down across the lake surface to the near hill. Observers were allowed to look up and down but discouraged from looking side to side. They normally lowered their heads to see the ball that served as a marker for the part of the hill they were to judge. Once the observer's gaze was oriented toward the ball and the observer was preparing to give a verbal estimate, a record of head orientation was recorded. Their verbal estimate was then recorded, the display went blank, and they were instructed to close their

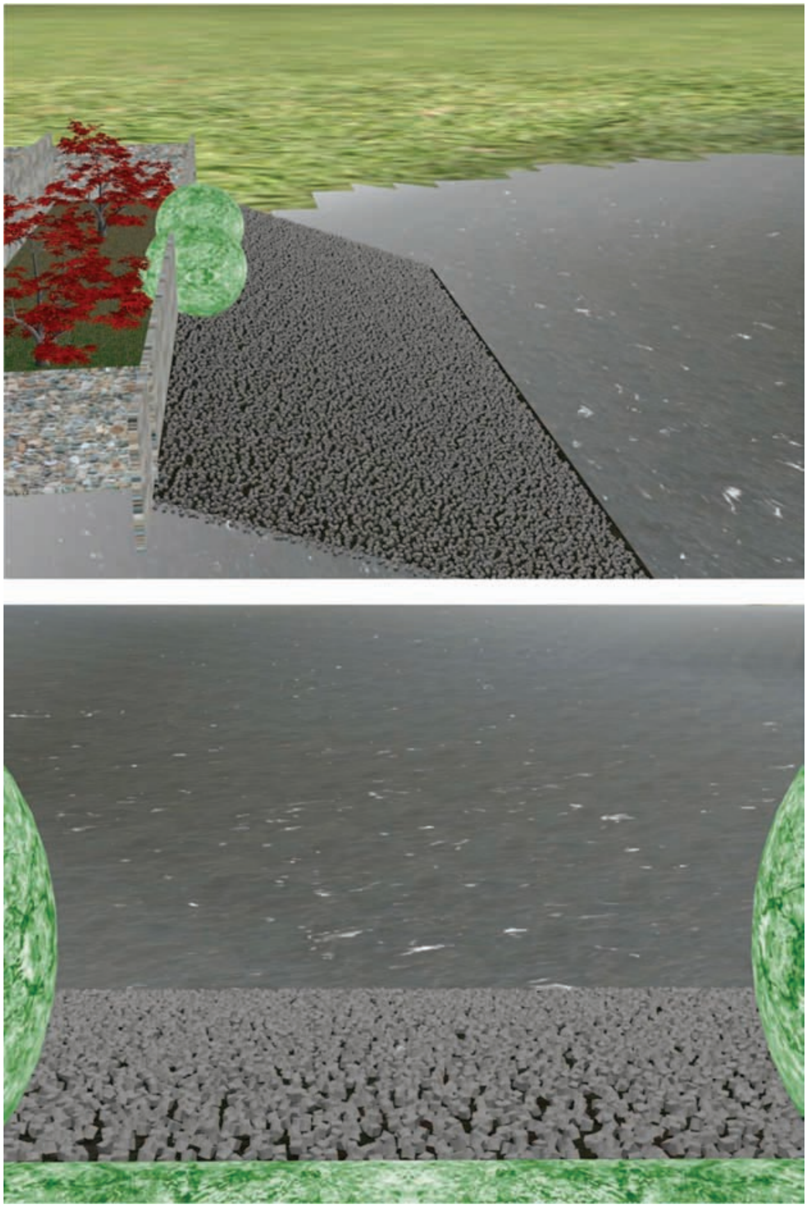

Figure 3. Virtual environment used for Experiment 2A. The top image shows overview of simulated scene. The bottom image shows view of down-hill slope from participant's point of view when standing back from edge of hill (Far condition). Curvilinear marble surfaces are visible at the bottom and sides of the frame.

eyes and hold their head upright again, as if looking straight ahead.

Results

Mean verbal estimates for each of the interpolated slopes are shown, by condition, in Figure 4. As in Experiment 1, for slopes of greater than about 16 degrees, there is a clear separation between the curves based on whether the viewing position was near to or far from the edge of the slope. Slopes appeared shallower when viewed from the edge.

For the main comparison between the Far and Near conditions, a mixed-effects model that included distance from the edge as a factor explained reliably more variance than one that did not, $\chi^2(1) = 49.2$, $p < .0001$. The model estimate of the effect was 7.7 degrees (95% CI = 2.0–13.4 degrees).

The effect was clearly not due to viewing distance per se because estimates in the Far-control condition were higher. A mixed-effects model including Far and Far-control conditions with Ball location as a factor explained reliably more variance than when Ball location was not included, $\chi^2(1) = 8.43$, $p = .0037$. The model estimate of the effect was 2.8 degrees (95% CI = 1.0 to 4.7 degrees).

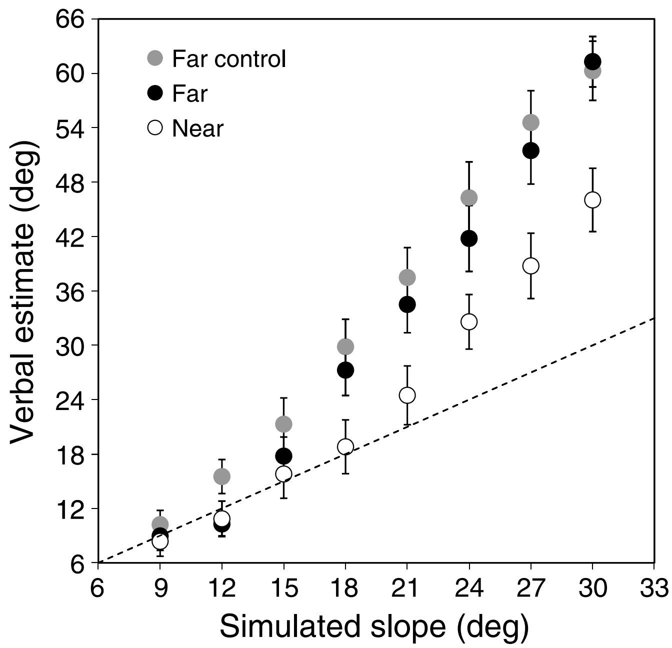

Figure 4. Results of Experiment 2A. Mean verbal estimates of simulated slopes viewed from the edge or from a meter back from the edge. Standard errors bars are shown.

It might be noted that the slope estimates for low slopes are less exaggerated than is typically reported for real hills; this may reflect a limitation of virtual representations. However, the clear separation between estimates at the edge and those from even 1 m back from the edge is consistent with our own observations in the real world.

Experiment 2B: Proprioception of head orientation—magnitude production

As a person steps back from the edge of a hill, one's gaze along the surface of the hill can approach parallelism with that surface. If gaze orientation were correctly perceived, such a circumstance should lead to more accurate (less exaggerated) perception of surface slant. Our results show the opposite, which suggests that direction of gaze is not perceived accurately. Our own observation outdoors was that our felt head orientation was exaggerated by a factor of about two, even with our eyes closed; this could contribute to misperception of gaze declination. In Experiment 2A, viewing the slope always required declination of the head. Misperception of head orientation (and thus gaze direction) could help account for the misperception of downhill surface orientation. To our knowledge, the proprioceptive misperception of head orientation has not previously been documented.

Methods

To quantify this exaggeration, we asked a new set of 13 participants to produce specific magnitudes of forward head pitch while wearing the HMD. The participants were blindfolded inside the HMD because we were interested in assessing non-visual proprioception of head orientation. We asked each participant to position the head at vertical between each trial and then to produce a specified head declination. In two randomly ordered blocks of five trials, declinations of 10, 20, 30, 45, and 60 degrees were requested. Head pitch data were recorded using the HiBall optical tracking system.

Results

The mean settings are shown in Figure 5. The produced declinations varied with a gain of only 0.5, suggesting that perceived head orientation increased with a gain of 2 relative to actual changes in head orientation.

The misperception of head orientation found here is not due to the weight of the HMD. Using a Vicon tracking system, we have measured comparable proprioceptive exaggerations in other participants wearing only a blindfold and a light hat (with a gain of about 0.4).

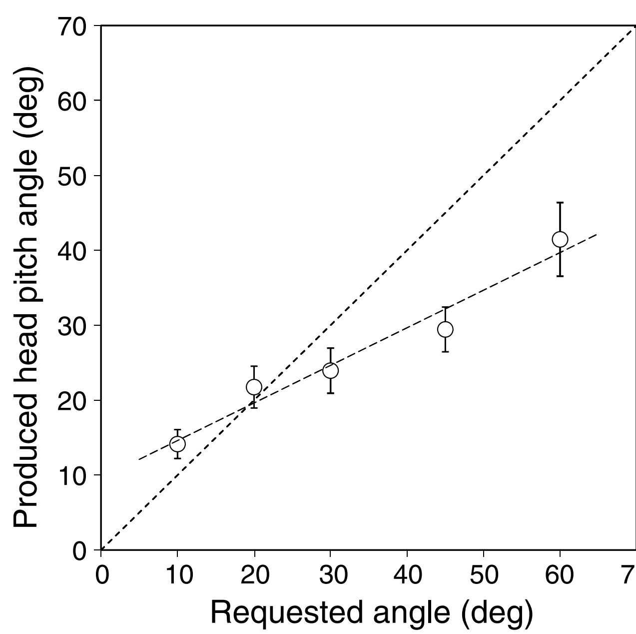

Figure 5. Results of Experiment 2B. Mean head orientations produced in response to requested angles along with least squares fit, which has a slope of 0.5. Standard error bars are shown.

A geometrical model of changes in downhill slope perception with point of view

Here we present a model of geographical slope perception based on three assumptions concerning the perceptual variables corresponding to the physical variables depicted in Figure 6.

1. Slope perception is an arithmetic function of perceived gaze orientation and perceived optical slant. Specifically, because geographical slope, α, is the difference between gaze declination, γ, and optical slant, β, as shown in Figure 6, perceived slope, α_p, is assumed to be the difference between the latent variables of perceived gaze orientation, γ_p, and perceived optical slant, β_p:

$$\alpha_p = \gamma_p - \beta_p. \quad (1)$$

2. Perceived gaze orientation, γ_p, may be a linear function of actual gaze orientation, γ, with a constant gain factor, k:

$$\gamma_p = k \cdot \gamma. \quad (2)$$

3. For any given surface, perceived optical slant, β_p, is a continuous function of actual optical slant, β.

$$\beta_p = f(\beta). \quad (3)$$

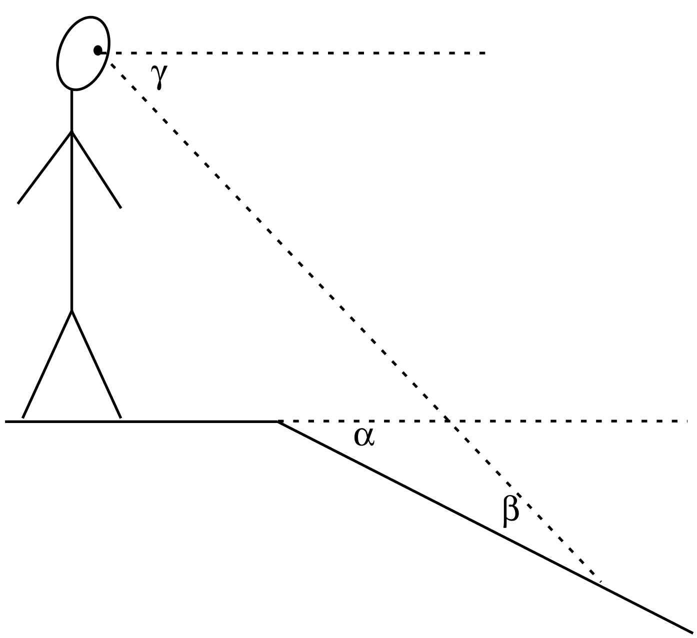

Figure 6. An illustration of the geometric relationship between downhill slope, α, gaze angle, γ, and optical slant, β, during downhill viewing, where $\alpha = \gamma - \beta$.

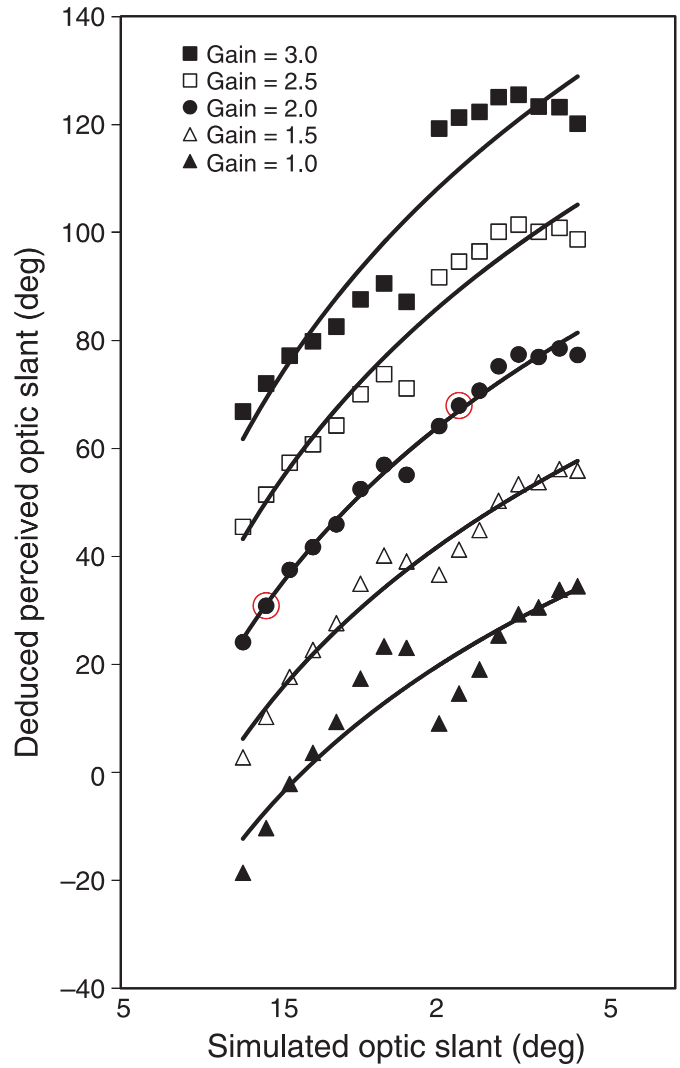

Figure 7. Results of Experiment 2A re-plotted as perceived optical slant derived from the model (Equation 4) by subtracting average slope estimates from the gaze declination to the simulated ball in the virtual scene, where gaze declination is multiplied by gains from 1 to 3. The break between the points to the left and right within each curve represents different viewing distances. The points circled in red represent model values for the 27-degree slope viewed from far (left) and near (right). Lines are logarithmic fits to the data. Note that optic slant is defined only from 0 to 180 (see Figure 6), following Sedgwick (1986).

If we apply Equations 2 and 3 to Equation 1, then we get the final model:

$$\alpha_p = k \cdot \gamma - f(\beta). \quad (4)$$

Early work on the perception of surface orientation focused primarily on perceived optical slant, the orientation of a surface relative to gaze (e.g., Flock, 1965; Gibson, 1950; Gruber & Clark, 1956; Perrone, 1982), rather than on geographical slant (Gibson & Cornsweet, 1952). Others have observed that errors in perceived gaze orientation may contribute to the misperception of surface orientation relative to gaze (e.g., Perrone, 1982), but such models have been concerned with the perception of optical slant rather than geographical slant—the orientation of a surface relative to horizontal (Gibson & Cornsweet, 1952) or have concerned optically induced misperceptions of the horizon (Bressan, Garlaschelli, & Barracano, 2003).

Taking the averaged slope estimation data in the Near and Far conditions of Experiment 2A, we can create a family of possible curves (Figure 7), based on Assumptions 1 and 2, that relate perceived optical slant (inferred from the data by subtracting judged slope from an estimate of perceived gaze angle) to simulated optical slant. With k set to 1.0, 1.5, 2.0, 2.5, and 3.0, we can see the resulting relationships between simulated optical slant and the perceived optical slant implied by Assumptions 1 and 2 in Figure 7. It is evident that Assumption 3 (i.e., that the relationship between simulated and perceived optical slant be a continuous function) is best met when the gain factor k is set to about 2.

Because we measured head orientation during Experiment 2A, we can also apply the gain factor directly to head orientation as if perceived gaze orientation were simply equivalent to perceived head orientation. This produces the same best fit with a proprioceptive gain of about 2.

Methodological considerations in studying failures of constancy

Granrud (2009) has shown that children who are able to articulate principles of constancy ("things look smaller when they are far away") appear to demonstrate much better size constancy than their age-mates. He argues that adults may also cognitively compensate for normal and pervasive failures of size constancy. When we ask people whether they think hills will look steeper from the edge or from back from the edge, we have found that most predict that they will look steeper from near the edge. Nonetheless, people may have more accurate implicit knowledge of how changes in location truly affect apparent surface slant.

When we attempted to develop a real-world version of the VR experiment, we first tried allowing our participants to reposition themselves from trial to trial. We discontinued the experiment when it became clear that no differences in verbal slope judgments were evident based on the manipulation of position, even though we changed the surface slant each trial. Although it was possible that people were dampening the effect by trying to "cooperate" with the wrong inferred hypothesis (Durgin et al., 2009), we suspected instead that they were able to implicitly compensate for the non-constancy that seemed evident to us. After all, adults do not usually notice failures of constancy. This may simply mean that, rather than having true constancy, they are able to predict the perceptual consequences of their actions as they move (e.g., Durgin, Gigone, & Scott, 2005), which is just as useful. Note that in Experiment 2A, participants did not move themselves, but the virtual world was simply repositioned around them.

The method used in Experiment 3A was developed to address our concern that cognitive strategies might mask the perceptual phenomenon we sought to expose.

Experiment 3A: Failures of geographical slant constancy with visual estimates of real surfaces

Whereas our model fit the data from Experiments 2A and 2B surprisingly well, the limited field of view of the HMD may have produced unusual viewing strategies. We therefore sought to use real surfaces to create a within-subject data set rich enough for modeling. In a companion experiment (Experiment 3B), we directly measured perceived declination of gaze.

In addition to the methodological considerations above, we sought to improve on Experiment 1 by controlling for retinal size and removing local relative orientation information. Although we regarded the between-subject design of Experiment 1 as a strength, a within-subject design was adopted here to facilitate modeling.

Methods

Participants were 25 students. Twenty-three had not been in any of the previous experiments. The remaining two were naïve about the experimental hypothesis.

Design

The main design of the experiment was to present three different surface orientations, downhill surfaces of 16, 24, and 32 degrees, at two different viewing displacements (0.6 and 1.6 m). In order to control for retinal size, large (~80 cm across) and small (~40 cm) surfaces were presented in each cell of the main design. To camouflage the design, these 12 experimental trials were randomly intermixed with 12 filler trials (some with large, some with small surfaces) that included three additional viewing distances (1.1, 2.1, and 2.6 m) and six additional surface orientations (8, 12, 20, 28, 36, and 40 degrees) so as to overload memory for each variable. (As a manipulation check, a questionnaire at the end of the experiment assessed participants' beliefs about the experiment, including the experimental design.)

To avoid the vagaries of verbal responses, the estimate was given by orienting a line within a circular aperture on a horizontal LCD screen (Figure 8). The anti-aliased line was constrained to represent an angle between −10 and 100 degrees of (downhill slope) to the nearest tenth of a degree. Its initial orientation was selected from a uniform random distribution in this range on each trial.

Apparatus and procedure

The sloped surfaces were 27 planar wooden surfaces (9 large and 18 small), cut into irregular shapes and protected with clear lacquer. A few are pictured in Figure 8 along with the mechanical apparatus used to precisely orient the surfaces. One experimenter repositioned the orientation of the device and placed a new wooden surface on it for each trial. The true orientation of the surface was measured with an inclinometer and recorded.

The stand on which the surface was mounted was surrounded by crumpled black felt cloth and the area beyond the surface was built up into an irregular mountain covered with black felt so that the immediate visual context for viewing the wooden surfaces provided no horizontal and vertical references. The entire apparatus was aligned along the diagonal of a large room and the more distal surroundings were visible in the periphery.

Participants viewed the surfaces from a wheeled platform (Figure 9) that could be smoothly repositioned between trials along a track. The floor of the platform was 29 cm below the center of the wooden surfaces to be judged. The platform was repositioned between trials by a second experimenter while the participant turned their head to the side. Each repositioning of the platform involved several movements back and forth to maximize perceptual uncertainty about the amount of final displacement.

The judgment line was presented on a horizontal LCD screen at about waist level on the platform. The anti-aliased white line was 10 cm long and 2 mm wide, drawn within a circular black window against a gray surround. Participants controlled the orientation of the judgment line by sagittal movements along a touch pad and pressed a key to indicate a match. Participants were trained in the use of the touch pad to adjust the line prior to the experiment.

Randomizations were prepared for each participant in advance. Before each trial, one experimenter moved the cart into position, while the other placed and measured the stimulus surface. The participants turned their heads to one side and closed their eyes during this time. When signaled to do so, they opened their eyes and looked at the wooden surface. They then made adjustments to the oriented line, looking back and forth between the line and the board. The completion of 24 trials took 30–40 minutes. With participant consent, the entire procedure was videotaped.

At the conclusion of the main experiment, a computerized survey asked participants their beliefs about the experiment: what they thought the experiment was about, how difficult the task was, how comfortable they had been, how many different board sizes, viewing distances, and angles there had been, what strategies, if any, they had used, and what they thought our hypothesis was.

Figure 8. Sample surfaces, the slope apparatus (upper right), and the estimation stimulus (lower right) used in Experiment 3A.

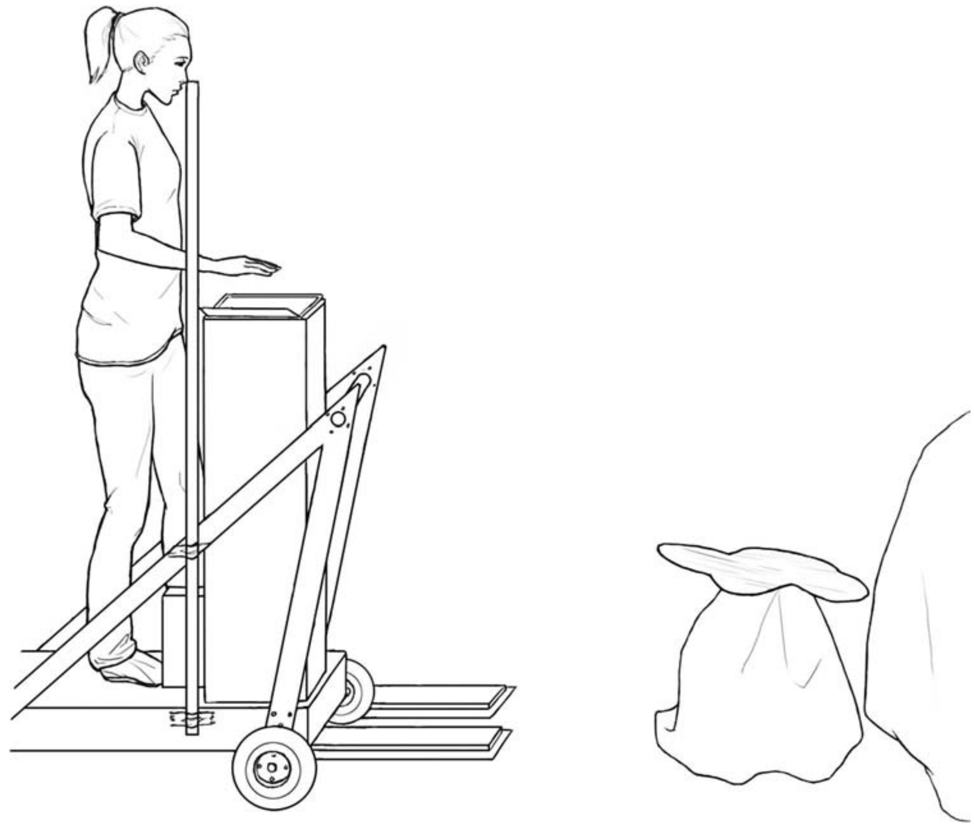

Figure 9. Moving viewing platform used for Experiment 3A.

Results

Survey data

The mean responses about the design were that there were 3.1 board sizes and 4.3 distances. The modal response was "4 to 6" different slants (five ranges were queried). Apart from underestimating the number of distinct slants, there was fairly good, but imperfect, cognitive awareness of the design. (Board size was ambiguous because the larger irregular shapes varied substantially in area.)

Most people (20 of 24 who completed the survey) guessed that our hypotheses concerned effects of distance, but only one person attributed to us the correct sign of our hypothesis, whereas seven suggested the hypothesis was that slopes would appear *steeper* from the nearer position. Four thought the hypothesis was that slopes were harder to estimate from very close and eight described no specific prediction. Among the eight who remained uncommitted about our hypothesis, two had described strategies involving compensating for the flatter appearance of the slopes viewed from close up!

Comparisons of individual data patterns and beliefs about our hypotheses suggested that a few participants may have been "compliant" with the wrong hypothesis (see Durgin et al., 2009), but there seemed no clear or urgent grounds for eliminating any data.

Orientation match data

The match data were analyzed with a 2 (Viewing Distance) × 2 (Size) × 3 (Slant) repeated-measures ANOVA. As expected, there was a reliable effect of Viewing Distance, $F(1,287) = 21.9$, $p < .0001$, in the predicted direction, as shown in Figure 10. The average difference was 6.3 degrees.

There was also a reliable effect of board size, $F(1,287) = 5.25$, $p = .0240$, such that large boards were judged about 3 degrees steeper than small boards. Because larger boards project larger retinal images, but farther boards project smaller retinal images, the direction of the size effect suggests that retinal size does not explain the effect of Viewing Distance. Neither Size nor Viewing Distance interacted reliably with Slant.

Modeling the data

Because there were effects of board size, we modeled the data for each board size separately. Figure 11 shows the behavior of the model (Equation 4) for data from the large boards with various estimates of the gain factor k. Perceived optic slant could be fit as a linear function of true optic slant for these data. We derived a best-fitting value of the gain factor for perceived direction of gaze by fitting the data iteratively and maximizing R^2. For the large

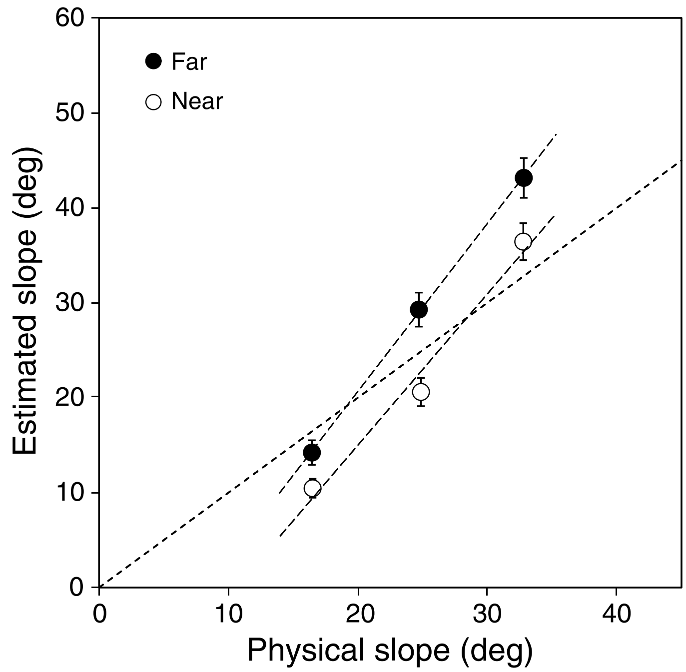

Figure 10. Results of Experiment 3A. Effect of viewing distance on orientation matches for real downhill surfaces. Standard errors with respect to within-subject differences are shown.

boards, the best-fitting gain factor was 1.54. For the small boards it was 1.39. The overall best gain factor was 1.50.

Observational data concerning head orientation

Our task required participants to visually orient with respect to two locations on each trial: the stimulus surface and the response screen. How much did they use their head to orient toward the stimulus?

Because we had video records of Experiment 3A (for all but two participants), a video frame in which the participant was viewing the surface was selected for each experimental trial. An image of the participant's head was cropped from the frame and enlarged. Two lab members unfamiliar with the design and one of the authors made estimates of head orientation using a custom program that allowed them to draw a line onto the photos. They attempted to draw the line parallel to the major vertical axis of the head and rated their confidence in their estimate. Orientation values were corrected for projective distortions of geometry in the camera view based on a calibration image of known orientations. For three videos, confidence ratings were consistently low, leaving useable the data from 20 participants. Mean orientation estimates were used.

Although the average gaze angle for the near position was 64 degrees and the average gaze angle for the far position was 38 degrees, participants were surprisingly consistent in orienting their heads. Heads were declined reliably further when viewing the near surfaces (34.3 degrees) than the far surfaces (31.8 degrees), $t(19) = 2.90$, $p = .0091$. But this represents a difference of only 2.5 degrees. Only two participants changed average head declination by more than 5 degrees. This suggests that participants may have sought to stabilize their head orientation in order to maintain a fixed frame of reference.

When measured head orientation was added to the model of perceived optic slant (by splitting perceived gaze declination into a head declination factor with a gain of 2.0 and an eye-in-head declination factor, with unknown gain), the best-fitting model still ascribed an overall gain of 1.46 to the eye-in-head component of gaze.

Discussion

As in Experiments 1 and 2A, a failure of constancy was found such that surfaces appeared shallower when viewed

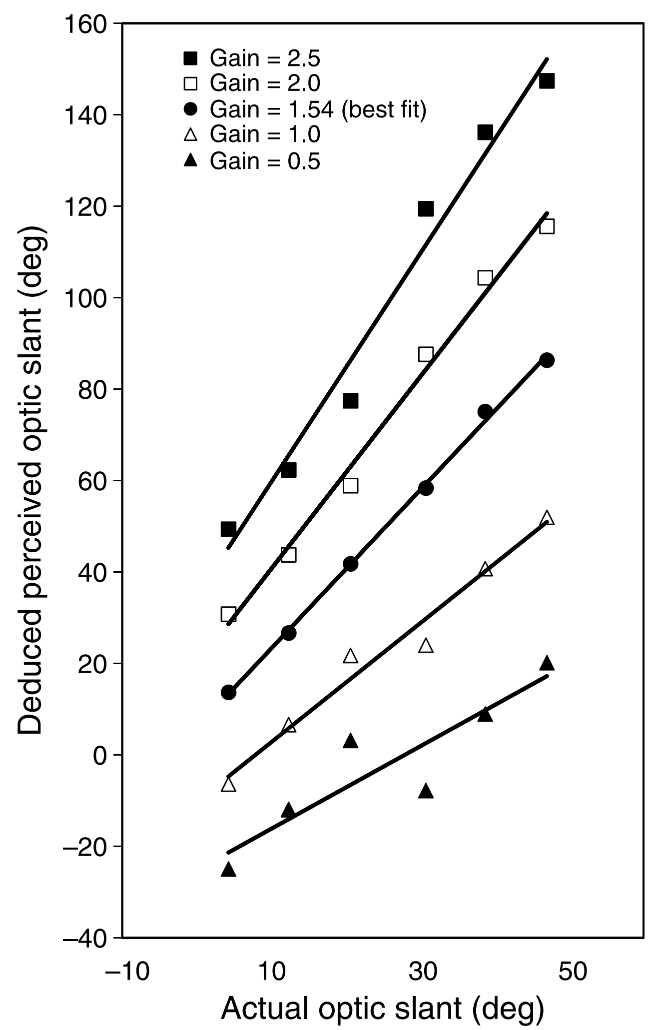

Figure 11. Model estimates of perceived optic slant as a function of true optic slant and various gains by which declination of gaze may be misestimated. Data are shown for the large surfaces in Experiment 3A. The best-fitting gain factor was 1.54. The slope of the regression line was 1.73.

with gaze farther declined. Although no participants mentioned trying to keep head orientation stable when describing their strategies, this may have been an implicit strategy many adopted to stabilize the frame of reference used for their judgments.

There were many differences between the VR set-up of Experiment 2A and the present experiment. Among these was that the field of view of the HMD in Experiment 2A limited eye-in-head declination to 15 degrees or less, so the head had to be declined to orient gaze. Whereas the gain factor derived from proprioception of head orientation in Experiment 2B provided a good fit for the data from VR of Experiment 2A, when applied to direction of gaze, a lower gain factor evidently fits the present data better, where participants had greater control over the amount of head declination.

Experiment 3B: Visual and verbal estimates of gaze declination

In unrestricted viewing conditions, gaze can be declined by a combination of head declination relative to gravity and eye declination relative to the head. Here we sought to measure perceived declination of gaze directly in a natural environment.

To avoid cognitive artifacts that might arise from participants analyzing scene geometry associated with the ground plane, we had participants view outdoor objects through windows.

Methods

The participants were eight undergraduate students.

Stimuli

We picked or positioned five landmarks to be viewed from five different windows in an upper floor of our building. For three of the scenes, small colored balls were anchored on the ground. The remaining two scenes used landmarks suspended in the air (a spherical light, a bend in a tree branch). The labeled views from the window are shown in Figure 12. The angles of the lines of sight from the windows for an average-height standing observer were 8.5, 15.8, 25.4, 34.7, and 43.3 degrees.

Procedure

Participants were shown how to adjust the orientation of a line on the screen of a laptop computer using the touch pad. They were led to each of the five views in a different random order and asked to set the orientation of the line to

Figure 12. Views of the scenes used for Experiment 3B.

depict the orientation of their line of sight to the target object. (In the case of the tree branch, a marked photograph was used to clarify which part of the branch was to be judged.) After they made their visual adjustment for all five scenes, they were led to the five windows again in a different random order, and this time they were asked to provide a verbal estimate of the angle between the line of sight and the horizontal.

Results

The results for both measures are plotted in Figure 13. The measured gain was 1.31 for the visual matches and 1.51 for the verbal matches. The one pair of estimates that seem quite low relative to the general trend (for 25.4 degrees) corresponds to the view in which there are strong rectilinear perspective cues specifying the ground plane (which was actually sloped slightly downward).

Discussion

In the absence of a ground plane at the feet, gaze declination is overestimated. The measured gain for perceived declination of gaze provides a good approximation to the best-fit gain estimated from the data of Experiment 3A using a simple geometric model of downhill slope perception.

We have not sought to exhaustively isolate eye-in-head proprioception but rather to demonstrate that estimates of gaze declination in real environments are distorted by multiplicative factors sufficient to account for the failures of constancy we have documented for downhill slope perception. Such perceptual error need not translate into errors of action (Durgin, 2009). Indeed, when measures of action have been used (e.g., walking to a previewed target), people seem to be using accurate estimates of gaze declination (Ooi, Wu, & He, 2001), but this may reflect calibrated action rather than accurate perceptual experience.

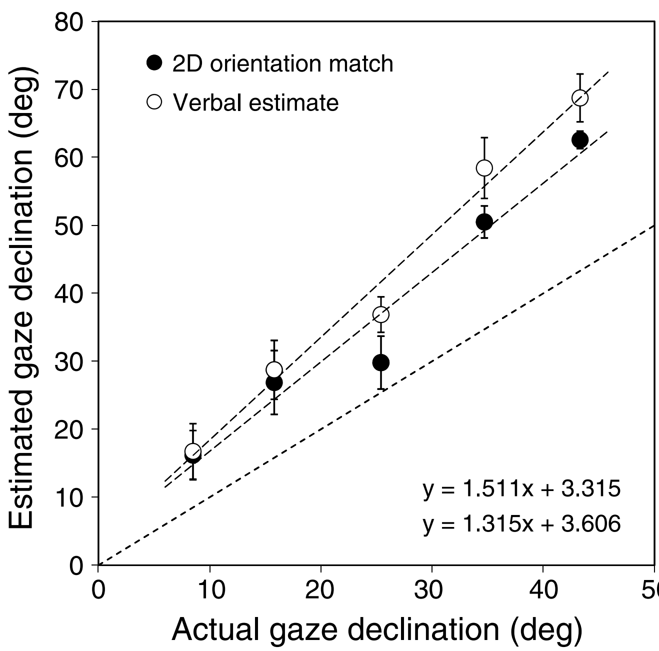

Figure 13. Results of Experiment 3B. Visual and verbal estimates of declination of line of sight from horizontal as a function of true gaze declination.

General discussion

We have measured striking failures of slope constancy for downhill slopes. We first observed the phenomenon with large-scale outdoor surfaces and have measured it for large-scale virtual surfaces using high-fidelity VR (Experiment 2A) and for small wooden surfaces (Experiments 1 and 3A). We have modeled the effect using both verbal judgments and non-verbal measures and found in both cases that the data fit a geometric model that includes a multiplicative bias in perceived direction of gaze.

Consistent with this perspective, we have measured a striking bias in non-visual head orientation proprioception with a gain of about 2 and shown that perceived declination of gaze is overestimated by a factor of nearly 1.5 under naturalistic viewing conditions selected to minimize ground plane information.

For our large-scale VR experiment using verbal reports, the relationship between simulated optic slant and perceived optic slant appears to be approximately logarithmic, whereas for the near, small real slopes used in Experiment 3A, the relationship measured with a visual angle was linear. In both contexts, however, a simple gain factor applied to direction of gaze was sufficient to make the relationship between true and perceived optic slant continuous.

Theories of geographical slant perception

Gibson and Cornsweet (1952) noted that far surfaces appear more frontal than near surfaces, suggesting a frontal tendency in slope perception, and Bridgeman and Hoover (2008) have documented this using verbal measures and a proprioceptive measure (holding out the arm). Bridgeman and Hoover suggest that perceptual errors at long distances are, like the moon illusion, inconsequential.

Frontal tendency

Frontal tendency was evident in our virtual reality study in the form of the logarithmic functions of optical slant. That is, when one looks down upon a downhill surface from its edge, frontal tendency will tend to make the surface appear shallower (more frontal to gaze), whereas the greatest overestimation of slope seems to occur when gaze is nearly parallel to the surface of the hill so that optic slant approaches zero. Although viewing distance may also play a role, our data clearly suggest that direction of gaze is an important factor and it seems to be systematically misperceived. In the studies of Bridgeman and Hoover (2008), viewing distance is confounded with viewing direction. Experiment 2A showed that viewing direction, rather than distance, was a crucial factor. Future work may help to disentangle these two factors for uphill slopes.

Behavioral potential

Whereas we have proposed that geometrical considerations are powerful enough to account for our newly observed failure of downhill constancy, Proffitt et al. (1995) have reported that slopes look steeper from the top than from the bottom and suggested that this is due to the greater biomechanical difficulty associated with descending a slope. However, their empirical claim must now be qualified by the observation that estimates of downhill slope depend a great deal on where one stands, which seems inconsistent with the emphasis on difficulty of descent. Similarly, our observations appear inconsistent with a recent report (Stefanucci, Proffitt, Clore, & Parekh, 2008) that fear influences the perception of downhill slopes—a finding that has also been questioned on methodological grounds (Durgin et al., 2009). The observation that steep hills appear shallower when standing at the edge seems to us inconsistent with the idea that fear plays an important role in the perceptual evaluation of slope (though it may play a role in judgment processes).

Coding efficiency

Perceptual coding processes may tend to overestimate geographical slant in order to maximize coding efficiency (Durgin, 2009). For example, our model implies that perceived optical slant in Experiments 3A and 3B varied as a linear function of true optical slant, but with a gain of about 1.7 (Figure 11). Indeed, such a gain is indicated simply by our data considered at each viewing distance separately.

We doubt that direct measurement of optical slant will always produce distortions of this magnitude, but note that for the range of optical slants used here (5–50 degrees, measured from the line of sight), the uphill slope perception literature may be interpreted as suggesting a similar scale expansion (e.g., Proffitt, 2006).

The relationship of perceived optical slant to actual (or simulated) optical slant presumably must depend on many factors related to the quality of the visual information available (e.g., texture gradients, stereopsis, etc.). Details of these considerations go beyond the scope of this paper but are reviewed by Flock (1965), Perrone (1982), and Sedgwick (1986). Issues of coding have also been discussed by Berends, Liu, and Schor (2005) and by Kinsella-Shaw, Shaw, and Turvey (1992).

Our observation has been that perceptions of both gaze direction and optical slant are scaled in a manner that exaggerates both, and we suggest that such biases may reflect coding efficiencies (sensitivity gains) associated with scale expansion (Durgin, 2009). Note that, as shown in Figure 14, optical slant and gaze declination are identical when looking at the ground plane. Thus, scanning a horizontal ground plane may tend to produce a calibration between the scaling of these two variables. The fact that both are exaggerated (and by similar amounts) may indicate that there are indeed coding gains

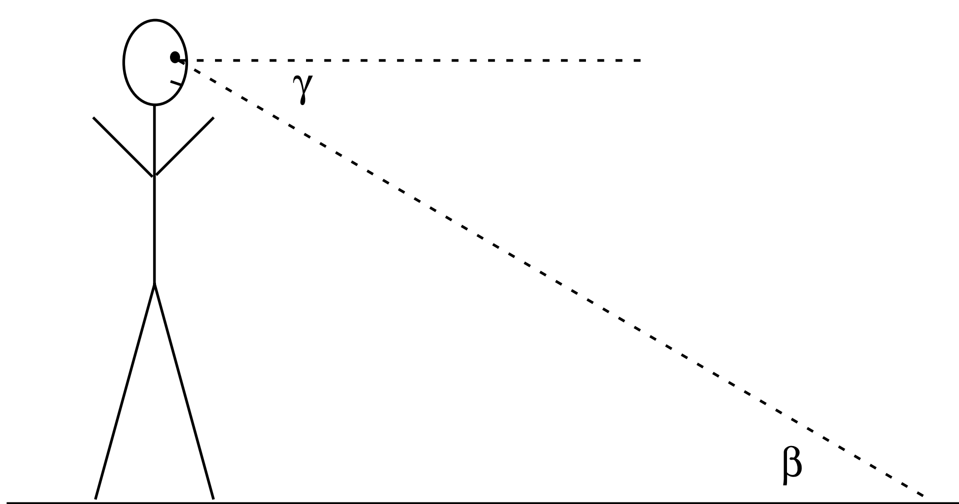

Figure 14. Along a horizontal ground plane, the angle of gaze declination is equal to the optical slant at the center of gaze. Equivalent misperceptions of both variables could still provide the appearance of constancy for the horizontal.

to be had by these scale expansions, even though they combine to produce (harmless) non-constancies when confronted with sloped surfaces. We speculate that the dual exaggerations of optical slant and direction of gaze may normally underlie our sensitivity to departures from the horizontal plane.

Conclusions

Hills look steeper when standing back from the edge and looking along them than when standing near the edge and looking down on their surface. We have shown that perception of gaze orientation itself and proprioception of the pitch of the head in particular are both overestimated. Simple geometric models of our data that take into account the misperception of gaze direction provide excellent fits to the data when the only free parameters in the model are measured empirically. In both sets of data we have modeled, the perception of optical slant is exaggerated in ways that are predicted by theories of coding efficiency. The misperception of gaze direction and the misestimation of optical slant may approximately offset each other in most visual contexts and normally provide coding advantages for detecting departures from the horizontal.

Acknowledgments

This research was supported by a Swarthmore College Faculty Research Grant to FHD. We thank Zijiang He for sharing his techniques for pincushion correction with us for the nVis HMD. Natasha Tonge assisted with conducting Experiment 3A. Carey Pietsch drew the illustrations.

Commercial relationships: none.
Corresponding author: Frank H. Durgin.
Email: fdurgin1@swarthmore.edu.
Address: 500 College Ave., Swarthmore, PA 19081, USA.

References

Berends, E. M., Liu, B., & Schor, C. M. (2005). Stereo-slant adaptation is high level and does not involve disparity coding. *Journal of Vision, 5*(1):7, 71–80, http://journalofvision.org/5/1/7/, doi:10.1167/5.1.7. [PubMed] [Article]

Bressan, P., Garlaschelli, L., & Barracano, M. (2003). Antigravity hills are visual illusions. *Psychological Science: A Journal of the American Psychological Society, 14,* 441–449. [PubMed]

Bridgeman, B., & Hoover, M. (2008). Processing spatial layout by perception and sensorimotor interaction. *Quarterly Journal of Experimental Psychology, 61,* 851–859. [PubMed]

Durgin, F. H. (2009). When walking makes perception better. *Current Directions in Psychological Science, 18,* 43–47.

Durgin, F. H., Baird, J. A., Greenberg, M., Russell, R., Shaughnessy, K., & Waymouth, S. (2009). Who is being deceived? The experimental demands of wearing a backpack. *Psychonomic Bulletin & Review, 16,* 964–969.

Durgin, F. H., Gigone, K., & Scott, R. (2005). Perception of visual speed while moving. *Journal of Experimental Psychology: Human Perception and Performance, 31,* 339–353. [PubMed]

Flock, H. R. (1965). Optical texture and linear perspective as stimuli for slant perception. *Psychological Review, 72,* 505–514. [PubMed]

Gibson, J. J. (1950). The perception of visual surfaces. *American Journal of Psychology, 63,* 367–384. [PubMed]

Gibson, J. J. (1966). *The senses considered as perceptual systems*. Boston: Houghton Mifflin.

Gibson, J. J., & Cornsweet, J. (1952). The perceived slant of visual surfaces-optical and geographical. *Journal of Experimental Psychology, 44,* 11–15. [PubMed]

Granrud, C. E. (2009). Development of size constancy in children: A test of the metacognitive theory. *Attention, Perception & Psychophysics, 71,* 644–654. [PubMed]

Gruber, H. E., & Clark, W. C. (1956). Perception of slanted surfaces. *Perceptual and Motor Skills, 6,* 97–106.

Kammann, R. (1967). Overestimation of vertical distance and slope and its role in moon illusion. *Perception & Psychophysics, 2,* 585–589.

Kinsella-Shaw, J. M. 1., Shaw, B., & Turvey, M. T. (1992). Perceiving "walk-on-able" slopes. *Ecological Psychology, 4,* 223–239.

Ooi, T. L., Wu, B., & He, Z. J. (2001). Distance determined by the angular declination below the horizon. *Nature, 414,* 197–200. [PubMed]

Ooi, T. L., Wu, B., & He, Z. J. (2006). Perceptual space in the dark affected by the intrinsic bias of the visual system. *Perception, 35,* 605–624. [PubMed]

Perrone, J. A. (1982). Visual slant underestimation: A general model. *Perception, 11,* 641–654. [PubMed]

Proffitt, D. R. (2006). Embodied perception and the economy of action. *Perspectives on Psychological Science, 1,* 110–122.

Proffitt, D. R., Bhalla, M., Gossweiler, R. & Midgett, J. (1995). Perceiving geographical slant. *Psychonomic Bulletin & Review, 2,* 409–428.

Sedgwick, H. A. (1986). Space perception. In K. R. Boff, L. Kaufman, & J. P. Thomas (Eds.), *Handbook of perception and human performance* (pp. 21.1–21.57). New York: John Wiley.

Stefanucci, J. K., Proffitt, D. R., Clore, G. L., & Parekh, N. (2008). Skating down a steeper slope: Fear influences the perception of geographical slant. *Perception, 37,* 321–323. [PubMed] [Article]