J Intell Robot Syst (2009) 54:595–612
DOI 10.1007/s10846-008-9281-y

Slope Perception from Monoscopic Field Images: Applications to Mobile Robot Navigation

Zhen Xiang · Geb W. Thomas · Kristopher M. Thornburg · Nathalie Cabrol · Edmond Grin · Robert C. Anderson

Received: 12 November 2007 / Accepted: 21 July 2008 / Published online: 30 August 2008

Abstract When remotely navigating a mobile robot, operators must estimate the slope of local terrain in order to avoid areas that are too steep to climb or that slope so steeply downward that the operator would lose control of the rover. Although many rovers are equipped with sensor systems to aid the operator in this task, it is sometimes necessary to estimate slopes from two-dimensional images, either when planning operations or when the operator wishes to monitor the results of a sensor system. This experiment compares the operator's estimates of the slope in Martian terrain with the actual slope determined from three-dimensional data. The ten participants overestimated the slope of the indicated regions by an average of 19° (SD 16°). An analytic model of the error, based on psychophysical analysis, accurately predicts the average magnitude of the errors. Implementation of this model eliminates an average amount of participant error. However, the large estimate variance within and between participants and images still poses a challenge for accurate slope estimation.

Keywords Slope perception · Slope estimation · Robotic teleoperation · Slope perception model · Image perception

Sheridan's [15] Supervisory Control Theory outlines how operators control most semiautonomous mobile robots or rovers. The control architecture explains that an operator controls a semiautonomous machine by specifying goals that the robot

Z. Xiang · G. W. Thomas · K. M. Thornburg (✉)
The University of Iowa, Iowa City, IA, USA
e-mail: kthornbu@engineering.uiowa.edu

N. Cabrol · E. Grin
NASA Ames/SETI, Moffett Field, CA, USA

R. C. Anderson
Jet Propulsion Laboratory, Pasadena, CA, USA

can execute on its own. The theory describes five roles the operator must fulfill: plan, teach, monitor, intervene, and learn. Crucial to safely controlling a rover, the operator must conduct the plan and monitor roles completely, consistently, and accurately.

Planning [15] requires the operator to understand the mission goals, the rover's current state, its capabilities and limitations, and its physical environment. The operator must then conceptualize a task sequence that the rover can understand and execute to achieve the mission goals. Ideally, the task sequence would also optimize certain mission goals such as completing the task in the minimum time, moving the furthest distance with the least energy, or minimizing risk to the rover.

Monitoring [15] requires the operator to continuously compare the current state of the rover with expectations, searching for failures or process variations. Effective monitoring requires the operator to accurately predict how the rover should perform the specified tasks, interpret the information coming back from the rover, and detect and track small variations between the predicted and actual states. The operator's understanding of the rover, its behavior, and the environment should be detailed enough to differentiate between significant and insignificant differences in expected and actual performance. Although both planning and monitoring require a precise understanding of the geometry of the rover's environment, people have trouble perceiving this geometry with the information that a rover generally provides [2, 4, 8].

A thorough understanding of the local geometry is required for the operator to perform the plan and the monitor roles. When planning, the operator determines favorable paths for navigating the rover around obstacles or dangerously steep slopes. The final mission plan must accommodate any environmental constraints. Monitoring requires an accurate understanding of the environment's geometry, aiding the operator in developing a precise expectation of the rover's behavior as it traverses the environment. Operators cannot sufficiently complete the plan and monitor roles because of difficulty perceiving environmental geometry as presented by the rover [2, 4, 8, 10]. In order to design a system to aid operators in completing their roles adequately and efficiently, we must first understand the difficulty experienced by operators.

Common difficulties operators have when navigating a rover include disorientation, overestimation of geographical feature size, problems with determining rover orientation, and problems with distance and slope estimations. In short, many operators describe the navigation of a rover as "looking through a soda straw" [17]. Some systems currently use 3D or stereoscopic displays; however, these systems can lead to operator disorientation [17]. Casper and Murphy [2] describe operators having difficulty trying to determine if the rover was in the upright position while van Erp [18] showed that operators consistently overestimate distances while navigating a vehicle from a video monitor. A major limiting factor in navigation is operator overestimation of height in monoscopic images [7]. Noticing this effect, Woods et al. [19] highlight examples in which operators cannot determine whether a rover can pass through an opening or over an obstacle. Proffitt et al. [13, 14] have shown that people are notoriously poor at determining the angle of an incline in virtual as well as real environments. This problem becomes compounded in the absence of perceptual cues humans are accustomed to in every day life, such as man-made objects, motion cues, structured texture, or even straight lines. Removing these cues from an image drastically reduces the amount of information an operator can gather from the image, usually resulting in a guess. From these examples, it is evident that rover operators have little understanding of the rover's geographical environment.

In an effort to increase the operator's understanding of the rover's geographical environment, Lewis et al. [8] developed a gravity-referenced view (GRV) system, which rotates the camera view of the rover along one axis with respect to gravity. This system allows the operator to visualize the current roll of the rover, improving the accuracy of determinations of the rover's ability to traverse the environment, therefore avoiding some rollover accidents. In the experiments conducted by Lewis et al. [8], the GRV system significantly reduces but does not eliminate the number of rollover accidents. This system improves the monitor role of the operator, making it slightly easier for operators to estimate the rover's current orientation and environmental position. The operator still requires assistance in making geographic determinations during the plan role.

To make the role of plan manageable for the operator, an aid must present the operator with information about the geographical environment ahead. This understanding relies on an unbiased perception of the terrain or slope over which the rover must navigate. People cannot naturally make accurate estimations of slope; therefore, the system must accurately present slope information to ensure satisfactory completion of the plan role. Effectively designing an interface to present this information requires an understanding of how operators currently perceive slope and how the perceived slope will be used in a navigational decision.

1 Background

With respect to the observer's viewpoint along the z-axis, a slope has two components: pitch and roll (Fig. 1). Pitch is defined as the surface's rotation away from vertical about the x-axis, which is equivalent to the angle between the y–z projection and the y-axis. Roll is the rotation from the vertical about the z-axis, which is equivalent to the angle between the x–y projection and the y-axis [13]. Most slopes in

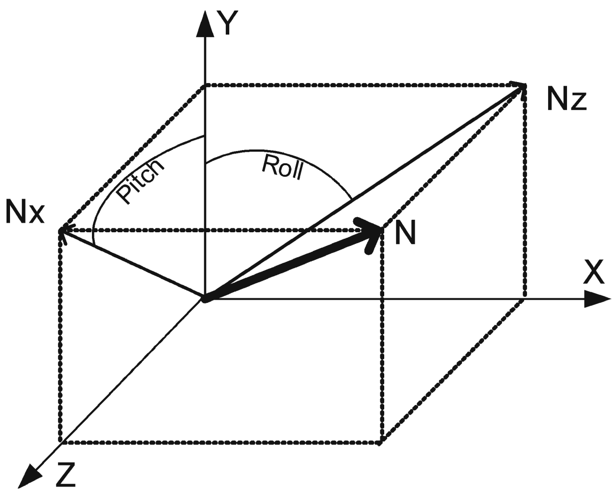

Fig. 1 Definition of pitch and roll. N is the normal of the terrain surface

a natural terrain have nonzero pitch and roll. Methods similar to the ideas employed by Lewis et al. [8] aid operators in determining roll but not pitch, so this paper exclusively considers the accurate perception of pitch.

Figure 2 demonstrates the various geometrical considerations involved in the estimation of local slopes with a rover. In this image, the rover's position is on a sloping terrain, so the rover's reference frame is at an angle, θ_r, with respect to gravity. In general, this angle could affect both the rover's pitch and roll but only rover pitch is considered here. Many rovers use a camera mounted on a pan-and-tilt unit, so that the camera can be directed towards different positions in the rover's environment. In Fig. 2, the camera aims downward at an angle t, which we refer to as the camera tilt. Horizontal displacement of the camera constitutes the camera's pan. The camera in Fig. 2 is imaging a region in front of the rover. This region has an overall slope tendency and a local region that has a separate slope tendency, defined by θ, which is the angle between the normal of the sloping surface and the camera's direct line of sight. The operator's task is to estimate the slope of the local region with respect to gravity. This slope estimate depends on the rover's angle with respect to gravity, the camera tilt, and the operator's estimate of θ. Generally, a gyrometer fixed to the rover's body measures angle with respect to gravity. Position encoders on the tilt and pan gimbal record the camera position.

It is not clear, however, how the operator incorporates the numerical estimates of robot tilt and camera tilt in his or her estimation of slope. For example, the operator may assume that the rover is level (i.e., $\theta_r = 0$) and estimate the camera pitch $(t + \theta_r)$ from the texture gradient of the general terrain in the scene. Such an assumption may be incorrect or prone to greater error than relying on the numerical values from the accelerometer and gimbal angle encoder. Sparse research covers how these values aid in improving slope estimations; however, we chose to explore the accuracy of slope estimates without these data to account for a worst-case scenario.

Many rover designers have circumvented operators' imprecise slope perception by incorporating hardware and software that directly estimate local terrain slopes with three-dimensional imaging systems, such as stereoscopic cameras or laser rangefinders. The operator can use this information to provide quantitative estimates of the

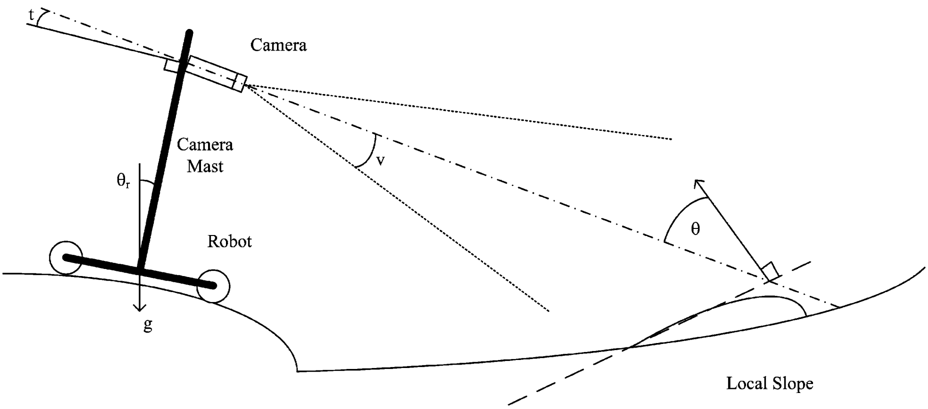

Fig. 2 Illustration of the task geometry involved in estimating a local terrain slope with an image taken by a rover equipped with a camera on a pan-and-tilt unit

slope or to generate a display of safe versus unsafe navigation regions. Because of the limits of sensor resolution, such solutions are generally only possible in the immediate vicinity of the rover. The operator must still estimate slope from images when onboard measurement systems are not available, when planning rover operations beyond the range of the onboard sensors, or when visually confirming the results of the automatically generated slope estimates. Work described here investigates the operator's ability to make these judgments without a three-dimensional map.

Existing perceptual models, developed with geometric situations much less complex than those investigated here, suggest that people's estimates of slope tend to bias in the direction towards the projection plane [5]. This bias is evident in laboratory studies of optical slant, in which research participants estimate the relative angle between their viewpoint and a stimulus surface (see [6, 16]). Research in geographical slant perception, in which research participants estimate the absolute slope of terrain when they are directly perceiving it or when they are looking at images of natural and virtual environments [3, 13], shows the same bias. This existing research suggests that observers misestimate the angle of regard or the projection plane of the visual scene. The research reported here considers the hypothesis that slope perception is inaccurate in the visual conditions provided by a mobile rover interface. More importantly, it is the first to demonstrate that this inaccuracy is large enough to limit effective rover operation.

There are several ways that an operator may improve performance over the conditions tested here. It is likely that training and experience with a rover would improve an operator's ability to perceive slope, so long as the training includes effective feedback. Without appropriate feedback, such as seeing the rover directly, an operator might form and maintain a false perception of scale, for example, imagining that the rover is larger than it is. Without seeing the rover approach and climb a steep slope, the operator might easily misconstrue how steep a slope the rover could climb. An operator might use familiar objects in the imagery to understand the scale of the environment. The operator might use landmarks to reconstruct the geometry of the terrain from different viewing positions and thus have an independent method for evaluating the environment's geometry. The presence of familiar objects or landmarks could create a perceptual feedback loop that might provide the feedback necessary to refine slope estimation with controlled experience with a rover. In unfamiliar environments where typical visual cues have been removed, performance may not improve beyond the unacceptable levels reported here. This research was taken on, in part, to determine if a support model can impact perception deficits.

Before developing a strategy to mitigate perceptual challenges, it is necessary to investigate the perceptual mechanisms that may be the source of the inaccuracy and to provide some baseline performance measure with which to understand the significance of the problem. For this, we consider two perceptual theories proposed by Perrone [10, 11].

1.1 Perrone's Two Models of Optical Slant Perception

Perrone [10] analyzed the texture pattern in visual images and concluded that a veridical perception of optical slant is possible if the observer can perceive direction and accurately estimate the length of the vector extending from the viewing position to perpendicularly intersect the extended plane of the slanted surface (line OA in

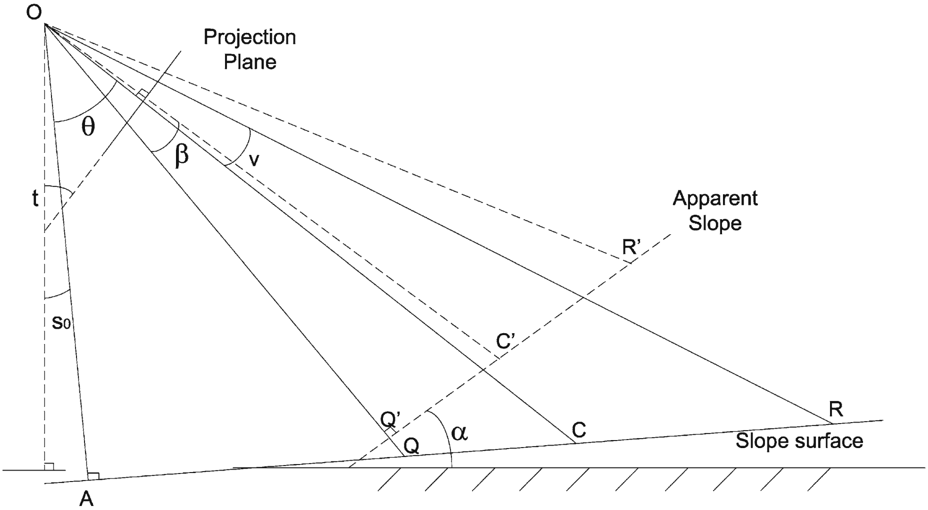

Fig. 3 Geometrical relationships for Perrone's Perpendicular Model. The observer at position *O* views the slanted surface, *DR*. The observer's vertical field of view is *ZV*. See text for definitions of other variables

Fig. 3). In practical situations, this critical location (position A in Fig. 3) is not directly perceivable by the viewer because of the limited field of view or frame. When the critical position is not observable, Perrone argues, observers form an assumption about the missing information. Perrone analyzes the consequences of two different assumptions that observers could make.

Perrone's first model (Fig. 3), which we refer to as Perrone's Perpendicular Model, proposes that observers erroneously take the slanted surface to be perpendicular to the ray along the bottom of the image (OQ). Perrone's rationale for this counterintuitive hypothesis is that the normal visual environment contains "an abundance of information about the direction of the perpendicular from the eye to the surface we are viewing." When the image frame removes this normally available information, the observer assumes that the point of intersection (A) is the bottom of the image (Q). With this assumption, the gradient is proportional to the angle of the optical slant, sin θ. In other words, Perrone proposed that, in the absence of the direct perception of the true perpendicular, observers assume that the vector OQ is perpendicular to the slope surface. Consequently, they perceive the slope to lie in the direction Q′ C′ R′ rather than QCR. As a result of this assumption, the perceived optical slant, β (<C′ OQ′), differs from the actual optical slant, θ (<COA). The two parameters are related by $\beta = \arccos(1 + \tan v \tan \theta)^{-1/2}$, where v is half of the visual angle vertically subtending the visible part of the slope surface. Despite the basis of Perrone's Perpendicular Model on optical slant, it extends to geographic slant. The next section considers the introduction of a variable for camera tilt, which aids in extending the model to geographic slant. In Perrone [10], this model is shown to be consistent with Gibson's [6] experiments with optical slant perception.

Perrone's second model, which we call Perrone's Reference Model, proposes that observers incorrectly assume that the line from observing point O to the bottom edge of the visible surface is parallel to the ground plane. The rationale for this hypothesis is that, in natural environments, the horizontal information is generally clear from the position of the visual horizon. When visual cues are impoverished, people are very poor at estimating the location of the horizontal direction, despite available posterial cues.

According to Perrone's Reference Model, the observer sees the surface in the correct relative position but misestimates the reference. If the observer takes the reference plane to be perpendicular to ground, the error in the slant estimate is equal to error in the reference plane position. Consequently, the perceived slant, ϕ, is given by: $\theta = \phi - v$, for slopes tilting away from the projection plane, as illustrated in Fig. 4.

Perrone compared the minimum absolute value estimate for the perpendicular and reference models using previously published optical slant estimates; the fit was within 3.25°. Later, Perrone and Wenderoth [12] followed by Perrone [11] proposed a more general model based on the idea that the direction from the observing point to the nearest part of the stimulus is perceived to be the true straight-ahead direction. The model may also apply to more general geometric cases, such as when the stimulus tilts forward relative to the visual line of sight. However, this revised model is essentially equivalent to Perrone's Reference Model, in which slopes generally tilt away from the projection plane, which is usually the case when judging potentially navigable terrain in natural environments. In addition, the later model considers the convergent angles of a regular texture pattern on the stimulus surface, reducing its applicability in natural environments. Perrone's models are the only models found to date of human slope perception published in peer-reviewed literature.

This paper seeks to measure slope perception error for rover navigation in natural, unstructured, desert-like environments to determine whether observers tend to overestimate slopes. The decision to study static, monoscopic, and naturalistic images containing no man-made objects or structured texture reflects the current needs for exploratory robots. By studying static and naturalistic image cues, the results of this study can be broadly applied as a worst-case scenario. This paper also seeks to determine which of Perrone's models may be adapted to this domain, given that

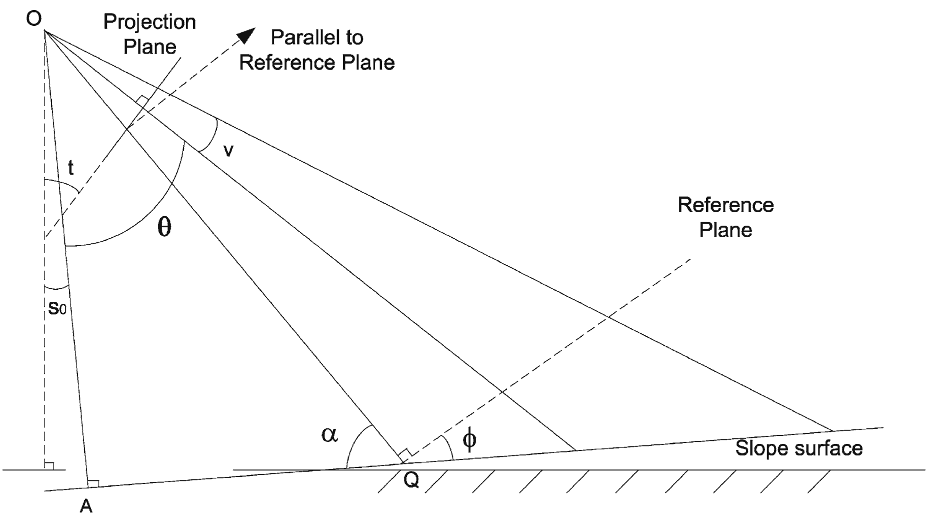

Fig. 4 Geometric relationships for Perrone's reference model

the problem of estimating slopes from images of natural environments considered here is somewhat more complex than the geometric situation considered by Perrone. Perrone's work relied on perceptual experiments in which a participant viewed a tilted planar surface through a round aperture. One complexity is that a natural visual scene may contain many regions with different slopes. A second complexity is that the observer cannot use the perception of gravity or isokinetic feedback to relate the optical slant to an absolute reference frame, a potentially large source of error. Instead, the observer must use the observed natural terrain to estimate the direction of gravity with respect to the viewed scene. This is analogous to estimating both the robot tilt and the camera tilt from the visual field of view. Previous perceptual experiments controlled or eliminated these factors [6, 10, 13, 16]. Consequently, it is not clear that the previous laboratory studies will directly apply to the problem of robot navigation.

1.2 Extending Perrone's Models to Rover Slope Estimation

To adapt Perrone's model to slope perception from field images, several additional complexities must be accounted for. First, both of Perrone's models assume that the observer views the physical stimulus directly. In this case, the perceived texture gradient is related to the observer's position relative to the textured surface. When viewing an image of a slope, the perceived texture gradient is a function of both the camera's position with respect to the surface and the observer's position with respect to the viewed image. This arrangement weakens the observer's ability to use his or her knowledge of body position relative to the surface and gravity. It also introduces the possibility of a distorted perception of the image if the observer's line of sight is not perpendicular to the image plane. Because we are interested in modeling the perceptual limitations of a rover operator when viewing a rover image, we do not attempt to restrict the viewer's head position or adjust the position of the image to match the camera tilt.

The second complexity arises from the fact that Perrone's models assume that a single sloping region fills the entire effective field of view, in which case the center of the slope and the center of the image are at the same position. When a region of slope is a subportion of the image, the center of the slope and the center of the image are different. To account for this, we calculate the geometric center of the slope and use this position. Consequently, there is a small offset term for the region offset added to the overall tilt of the image. Thus, we redefine t in Perrone's model to be t', as illustrated in Fig. 5.

A third complexity is the adjustment for half the vertical field of view, v, used by Perrone, which is also based on the assumption that the slope occupies the entire field of view. Instead, we use v', which is half of the vertical field of view of the sloping subregion of the image, which is also illustrated in Fig. 5.

The fourth complexity is related to the observer's need to estimate the absolute slope, rather than the relative slope modeled by Perrone. This is problematic because all the cues related to the direction of gravity are indirect. If the general tendency of the terrain is perpendicular to gravity, the global horizon and texture gradient provide information about the direction of gravity. If the global terrain is not perpendicular to gravity, however, the horizon and texture gradient will not reflect the direction of gravity, only the general trend of the landscape. For the sake of this

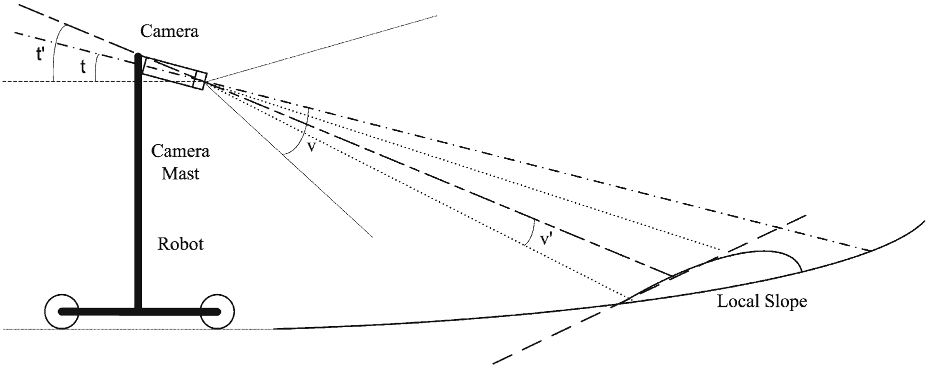

Fig. 5 Geometric considerations for extending the models

study, we choose to address this based on Perrone's two models and our experiment instruction. For the Reference Model, only the direct estimation of the tilt of local terrain is required. For the Perpendicular Model, the observer must perceive the direction of gravity in the image in order to estimate the camera tilt. We presume that the observers estimate the camera tilt from indirect cues, such as the texture gradient of the whole image and the position of the rocks and sediment, for example. It is not clear exactly what cues contribute to the perception of the camera tilt angle, but the images clearly provide some indication. This assertion leads to the assertion that the observer can also estimate the tilt of the camera, also through indirect cues.

If the observer can estimate the camera tilt, then he or she can estimate the absolute slope of a region by transforming the optical slant to global coordinates. For Perrone's Perpendicular Model (Fig. 3), the real slope QCR tilts up to Q′C′R′. In this situation, α is the slope perceived by observers and $\alpha = 90° - t' - v'$. This is the Modified Perpendicular Model. For Perrone's Reference Model (Fig. 4), OQ is assumed to be parallel to the ground; α is the perceived slope, and $\alpha = t' + s_0 + v'$, where S_o is as indicated in Fig. 3. This is the Modified Reference Model.

Our objective is to measure the accuracy of a novice observer's estimate of slope in conditions similar to those experienced by a rover operator and to determine whether an observer perceives slopes in an image of a natural environment in the manner predicted by the Modified Perpendicular Model or the Modified Reference Model. Specifically, we hypothesize that observers will perceive the slope of areas within a region according to either: (a) $90° - t' - v'$ or (b) $t' + s_0 + v'$.

2 Method

2.1 Participants

Ten participants (three females, seven males) aged between 18 and 60 years were recruited locally; eight were University of Iowa students. All the subjects have normal or corrected to normal vision. Novice operators were used in this study because training has been shown to not be very effective when navigating through an unstructured environment from a robot-mounted camera [1, 19].

2.2 Stimuli

The stimuli were taken from the Mars Pathfinder mission data archive [9], which includes stereoscopic image pairs taken by the Mars Pathfinder mission lander and a three-dimensional terrain model of the terrain within 2–15 m of the lander generated from these images (Jet Propulsion Laboratory [9]). The terrain model uses a set of over 90 images as texture maps. Each of these 256 × 248-pixel images subtends a 14.0° × 13.6° field of view with respect to the lander camera viewpoint. About 120 points in each image are tied to three-dimensional coordinates in the terrain model. The experimental stimuli consist of a subset of 15 natural environment images from the full dataset. Images were selected based on the following criteria: (1) selected images should not contain views of the lander; (2) selected images should display a range of slope angles at varying distances; and (3) the selected slopes should be planar and the uncertainty of the coefficients of regression for a plane passing through six selected points on the slope should determine pitch within 0.5°.

Estimates of the selected slopes in the image subset were calculated from the three-dimensional terrain model. To estimate a slope within an image, we first selected three points from the terrain model corresponding to locations in the sloping region in the image. An algorithm then considered all combinations of groups of six points in the three-dimensional data close to the three selected points. The set of six points having the largest R^2 value when fitted to a plane was taken to be the best group of points to represent the slope. This procedure resulted in six points associated with each selected slope, along with an estimate of the slope of the corresponding planar surface passing through the three-dimensional position of the points. All slopes in stimulus set have a confidence interval of ±0.5° at the 70% confidence level. A consequence of defining a mathematically unambiguous sloping

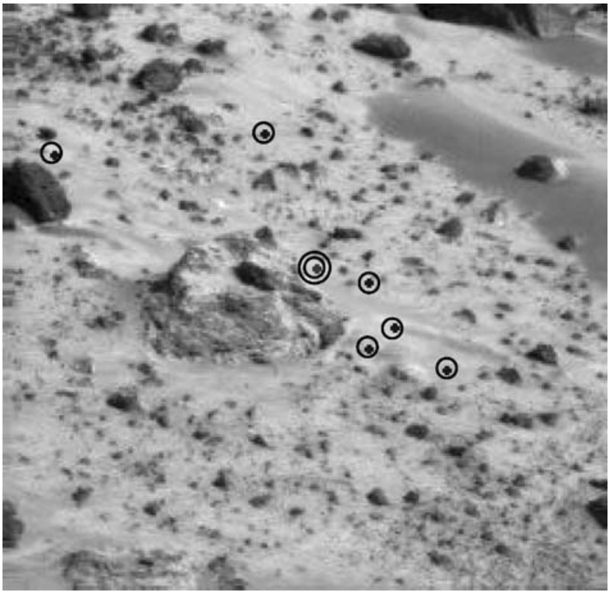

Fig. 6 A sample of a fully annotated stimulus. The *outlined circle* (which was *red*) indicates the geometric center of the slope; the other *circles* (which were *blue*) indicate the points used to indicate the local region of interest

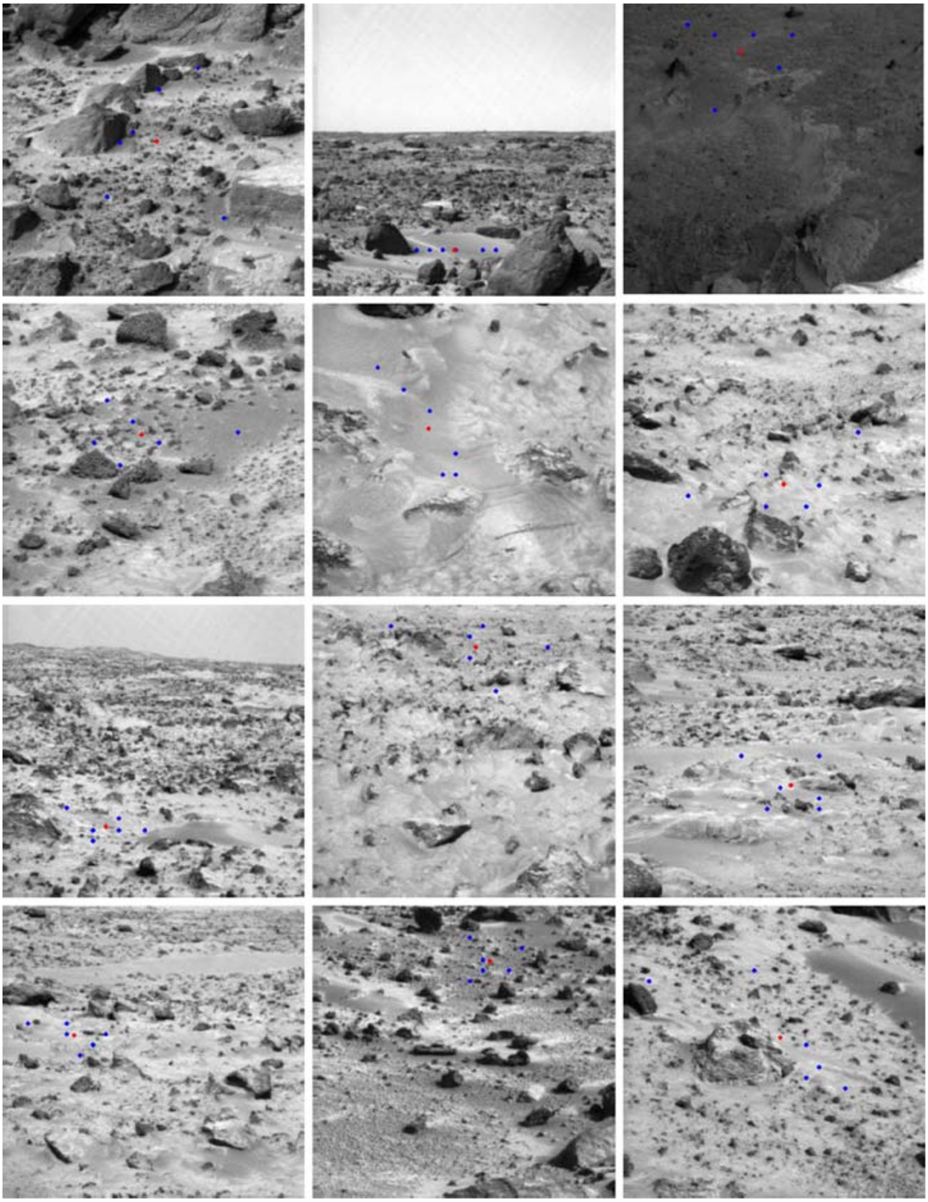

Fig. 7 Examples of stimuli images

region was that the points are not uniformly spaced, though it seems appropriate for natural environments with uncontrolled geometry.

Two images were prepared for each slope in the stimulus set. The first image was the original image annotated with a small, red, filled circle that indicates the geometric center of the points used to estimate the slope. The second image included the blue circles placed at the position of each of the sample points, plus the red

Fig. 8 The slope indicator. The circular top surface rotates about a pivot point below its center

circle placed at the center of the chosen points (Fig. 6). These annotations served to direct the participant's attention to the correct region of the image and specify exactly which points were considered to be included in the slope. Figure 7 shows a subset of images used in the experiment.

2.3 Apparatus

A laptop with a 15-in. screen displayed the images. Participants reported their slope estimates with a slope indicator, a flat circular plate covered with a random texture pattern mounted on a pan-and-tilt unit that tilted in both directions about the center of the plate (Fig. 8).

2.4 Experimental Design

Each subject estimated the slope of all 15 stimuli three times. Each stimuli set was presented in a separate, randomized block. For each trial, the pitch and roll angles of the slope indicator were recorded. The dependent variable is the estimated pitch angle, defined below.

2.5 Procedure

Participants adjusted their seat and the laptop screen angle to a comfortable working position. The distance from the participant's eye to the screen was measured, along with the height of the eye above the center of the screen. For each stimulus, first the image annotated with the location of all the points in the region was displayed, followed by the image with just the annotation of the geometric center of the slope. Participants could freely alternate between the two views, but their final slope estimate was made from the second annotated image. For each stimulus, the participants indicated their estimate of the local slope on the slope indicator so that the slope indicator surface would pass parallel to the plane formed by the slope.

Participants were instructed to imagine that their eyes were at the position of the camera and to indicate their perception of the absolute slope of the indicated region by adjusting the angle of the tilt indicator. They were also asked not to assume that the reference frame provided by the table top was parallel to the terrain but to exclusively rely on their perception of the top of the slope indicator to indicate their perception of the local slope.

3 Results

A total of 450 estimates were collected. Participants significantly overestimated pitch with a mean error of 19° (SD 16°, $t(449) = 24.68$, $p < 0.0005$). Figure 9 presents the pitch estimate error as a function of stimulus image.

For each image, the observed pitch errors (estimate minus actual) were compared with the errors predicted by the Modified Perpendicular and Reference Models. Figure 10 present the average and predicted errors of the Modified Perpendicular and Reference Models of each stimulus. A paired *t* test indicates that the errors

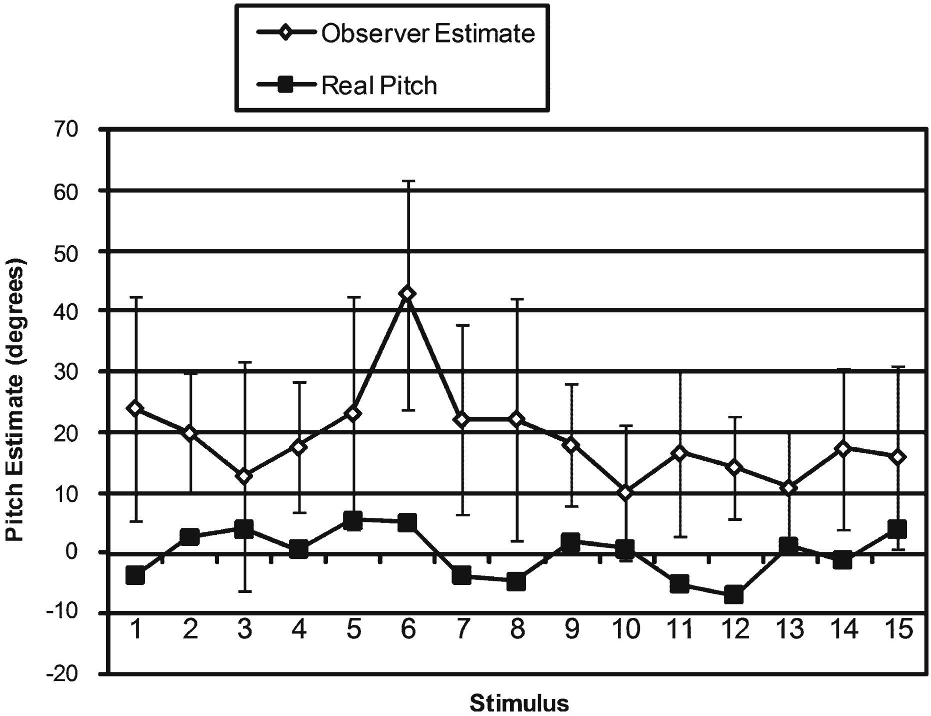

Fig. 9 Pitch estimates vs. stimulus image. The *error bars* indicate one standard deviation of the combined participant estimates

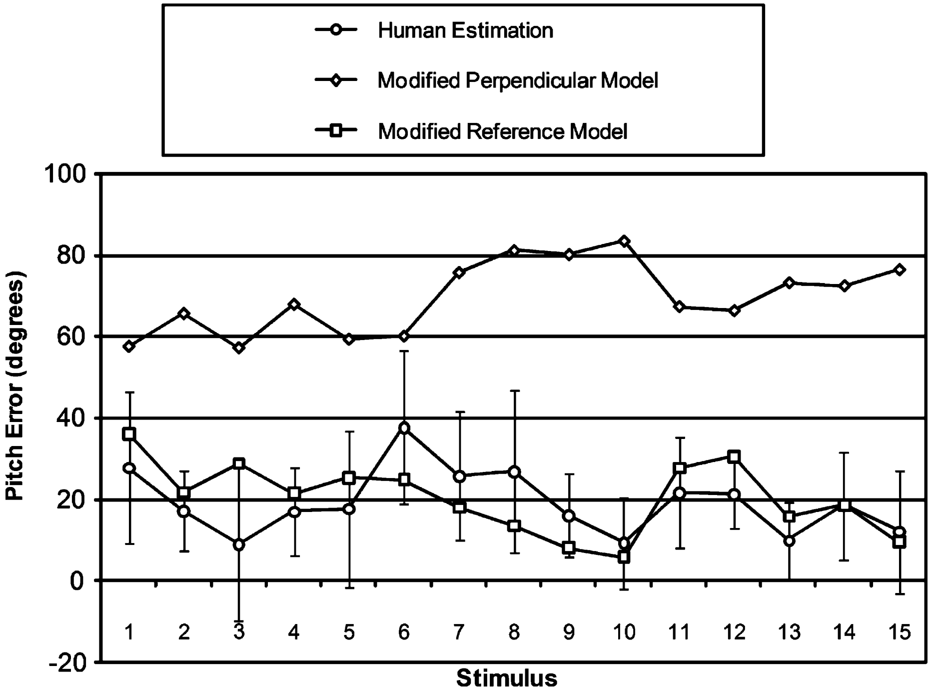

Fig. 10 The estimated error predicted by both models compared to the error observed in the experiment

predicted by the models are significantly different ($t(14) = 63.08$, $p < 0.0005$). The average error predicted by the Modified Reference Model is 47.9° less than those predicted by the Modified Perpendicular Model. The observer pitch estimate errors and the Modified Reference Model predictions are not significantly different ($t(449) = -1.5$, $p = 0.933$).

The average standard deviation of the subject pitch errors was 14°. The average standard deviation of each individual subject on each individual image was only 4°. A two-way analysis of variance of the pitch error as a function of subject, image, and subject-by-image interaction indicated that 26%, 22%, and 36% of the remaining variability is accounted for by the subjects, images, and subject-by-image interaction, respectively.

4 Discussion

The data suggest that absolute slope estimation from monoscopic images may pose a significant challenge for a rover operator. The average pitch overestimation of 19° is quite large compared to the size of slopes that may be safely traversed by rovers. If a rover operator overestimates the magnitude of an upward slope, he or she will likely navigate the rover to avoid the slope, which can potentially lead to choosing inefficient paths and missing opportunities to explore upslope areas that the rover is capable of navigating. The bias also suggests that operators may underestimate the magnitude of down slopes or estimate steep downward slopes to be level or even sloping upward. In this case, the operator might direct a rover down a potentially dangerous slope, leading to a loss of navigation control. This condition represents a serious potential rover hazard. The lack of structured texture and lack of reference to the rover are most likely the largest contributors to these gross overestimations or underestimations.

The data do not support the Modified Perpendicular Model, which proposes that observers assume that viewing direction to the lowest point in the projection plane is perpendicular to the ground between the bottom of the projection plane and the position of interest. The data do support the Modified Reference Model, which proposes that observers assume that the viewing direction to the lowest point in the projection plane is parallel to the ground. The between-stimulus differences between the observer's pitch estimate error and the Modified Reference Model suggest that other factors, such as the size of the slope or its position relative to the center of the image, may contribute to this slope estimation bias.

The Modified Reference Model provides an opportunity to mitigate this overestimation bias by correcting operator estimates. The model predicts pitch bias as a function of the camera pointing angle and the actual slope. During rover operations, the true camera angles are known, but the true terrain pitch is generally not known. If the model uses the estimated pitch from the rover operators as an approximation of the true pitch, it will predict a bias, which may then be subtracted from the operator bias.

Figure 11 presents the errors of observer estimates and the same estimates corrected by the Modified Reference Model, following the above procedure. A paired t test comparison of the errors of each observation with their corresponding corrected estimate indicates that the corrected estimates had significantly less error than the original estimates ($t(449) = 50.59$, $p < 0.0005$), with a 95% confidence interval (19.538°, 21.117°). The amount of correction is very close to the scale of original estimation error. A paired t test between the corrected estimates and the real slopes indicated that they were not significantly different ($t(449) = 1.5$, $p = 0.134$). The 95% confidence interval for the error of corrected estimates is (−0.373, 2.784). Consequently, when averaged across all the images, the error reduction is about 89–98%. For each single image, the model correction reduced the error by 40–99%, except for image 3, where the correction makes the estimation worse by 125.9%. No obvious feature of image 3 suggests why it would be estimated differently than the other images.

Further model refinements designed to reduce the between-stimulus differences may not be as important as reducing the variance in observer estimates. The standard deviation of 16° is too large for safe rover navigation, even if the estimates are unbiased, because the occasional large misestimation could lead to catastrophic mission failure. Compared to most rover navigation constraints, even an unbiased perception of slope with such a large uncertainty would be inappropriate for reliable navigation. Figure 12 highlights this fact by showing the 95% confidence interval for safe rover operation on slopes (dotted lines). This simple graph is exceptionally profound when showing that the negative slopes are consistently estimated to be very positive with large errors that frequently overlap the rover slope confidence. Consequently, the

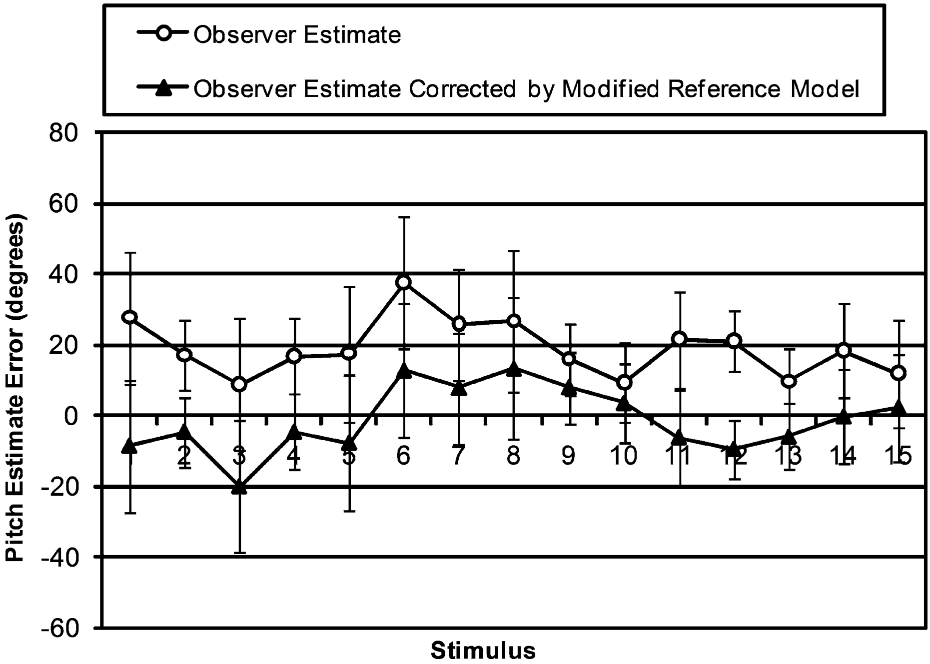

Fig. 11 Observer estimates corrected by the Modified Reference model

most pressing research question is to understand how to increase the precision with which rover operators estimate slope. By further refining the model based on geometric considerations, we could hope to eliminate as much as 22% of the variation

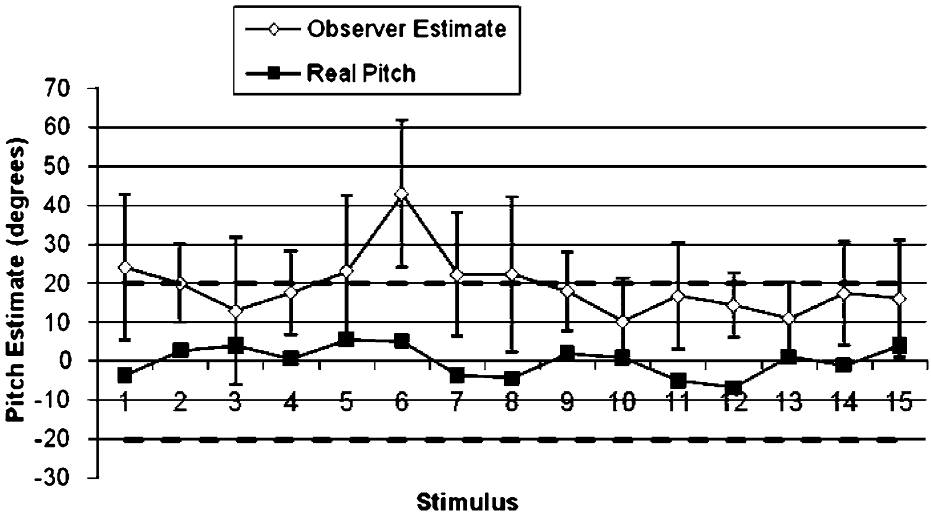

Fig. 12 Observer estimates for each stimulus image with confidence intervals for safe rover navigation (*dashed lines*)

in the errors associated with image differences. However, it may be more beneficial to study how operator pitch estimation strategies differ because 36% of the error variation is associated with the subject-by-image interaction. Without specialized training, our results suggest that participants are consistent in their estimates, as indicated by the 4° standard deviation of pitch error. With training, an operator may become even more consistent, but this has not been satisfactorily tested. An effective error range for slope estimation in rover operations would be on the order of 5% or a standard deviation of approximately 1° for a 20° slope. Increasing the estimation precision by a factor of approximately 4 would significantly reduce the rover's system complexity by eliminating some of the need to carry sensors to detect terrain slope and would increase the operator's ability to reliably supervise rover operations.

Future work in this area include implementing interactive visual aids overlaid on the images, such as a checkerboard pattern, to determine the impact on operator performance. By including an interactive visual tool, the operator may be able to refine the initial judgement when adjusting the visual tool. For example, the pan–tilt platform used in this experiment could control a checkerboard graphic overlaid on the image and adjusted by the model. However, this visual tool would still be subject to operators' perceptual biases. The same is true with sensor fusion techniques that add information available to the operator: the operators' perceptual biases will still affect the final judgement. Additional tools and techniques need to be investigated to provide more assistance for operators.

5 Conclusions

Observers overestimate the absolute pitch angles of a sloping region in the Mars Pathfinder dataset by 17.6° to 20.6° at the 95% confidence level. The bias may be corrected, on average, with a model that presumes that observers take the ray from the camera's focal point to the bottom of the locally sloping region to be parallel to the horizon. This model may be used to correct for observer slope pitch estimate biases during rover operations. However, the between-observer and within-observer pitch estimate variation still pose significant challenges for a rover operator attempting to estimate local slope from a monoscopic image. The data and literature suggest that humans cannot make accurate slope estimations in two- or three-dimensional images. In order for an operator to be able to independently assess the quality of three-dimensional terrain information presented by a rover's terrain-mapping subsystems or to plan rover operations beyond the range of such subsystems, it is necessary to refine appropriate technologies and training techniques. Devising a combination of training and technology to improve rover operator's ability to accurately perceive local slope would increase the operator's effectiveness in navigating the rover reliably and safely.

Acknowledgements This material is based upon work supported by NASA under award No: NEG5-11981. The authors would like to thank the members of the GROK Laboratory at the University of Iowa for their constructive insight and the research participants who participated in this study.

References

1. Burke, J., Murphy, R., Coovert, M., Riddle, D.: Moonlight in Miami: field study of human–robot interaction in the context of an urban search and rescue disaster response training exercise. Hum. Comput. Interact. **19**(1), 85–116 (2004)
2. Casper, J., Murphy, R.: Human–robot interactions during the robot-assisted urban search and rescue response at the World Trade Center. IEEE Trans. Syst. Man Cybern. Part B Cybern. **33**(3), 367–385 (2003)
3. Creem-Regehr, S.H., Gooch, A.A., Sahm, C.S., Thompson, W.B.: Perceiving virtual geographical slant: action influences perception. http://www.psych.utah.edu/~sc4002/pubs.htm (2003). Accessed 27 January 2004
4. Drury, J., Scholtz, J., Yanco, H.: Awareness in human–robot interactions. In: IEEE International Conference on Systems, Man and Cybernetics (2003)
5. Epstein, W., Park, J.: Gibson's psychophysical hypothesis. Psychol. Bull. **62**(3), 180–196 (1964)
6. Gibson, J.J.: The perception of visual surfaces. Am. J. Psychol. **63**, 367–384 (1950)
7. Kanduri, A.K., Thomas, G., Cabrol, N., Grin, E., Anderson, R.C.: The (in)accuracy of novice rover operators' perception of obstacle height from monoscopic images. IEEE Trans. Syst. Man Cybern. Part A **35**(4), 505–512 (2005)
8. Lewis, M., Wang, J., Hughes, S., Liu X.: Experiments with attitude: attitude displays for teleoperation. In: IEEE International Conference on Systems, Man and Cybernetics (2003)
9. NASA: Planetary Data System Database, Mars Pathfinder IMP imager, Presidential Panorama. http://stardev.jpl.nasa.gov/pds/index.jsp (2003). Accessed 11 November 2003
10. Perrone, J.A.: Slant underestimation: a model based on the size of the viewing aperture. Perception **9**, 285–302 (1980)
11. Perrone, J.A.: Visual slant underestimation: a general model. Perception **11**, 641–654 (1982)
12. Perrone, J.A., Wenderoth, P.M.: Visual slant underestimation. In: Ellis, S.R. (ed.) Pictorial Communication in Virtual and Real Environments, pp. 496–503. Taylor & Francis, London (1981)
13. Proffitt, D.R., Bhalla, M., Gossweiler, R., Midgett, J.: Perceiving geographical slant. Psychon. Bull. Rev. **2**(4), 409–428 (1995)
14. Proffitt, D.R., Creem, S.H., Zosh, W.D.: Seeing mountains in mole hills: geographical-slant perception. Psychol. Sci. **12**(5), 418–423 (2001)
15. Sheridan, T.: Telerobotics, Automation, and Human Supervisory Control. MIT Press, Cambridge (1992)
16. Smith, A.H.: Outline convergence versus closure in the perception of slant. Percept. Mot. Skills **9**, 259–266 (1959)
17. Steinfeld, A.: Interface lessons for fully and semi-autonomous mobile robots. IEEE Int. Conf. Robot. Autom. **3**, 2752–2757 (2004)
18. van Erp, J.: Trade-offs between spatial and temporal resolution in driving unmanned ground vehicles. In: Proceedings of the Human Factors and Ergonomics Society 42nd Annual Meeting, pp. 1550–1554 (1998)
19. Woods, D.D., Tittle, J., Feil, M., Roesler, A.: Envisioning human–robot coordination in future operations. IEEE Trans. Syst. Man Cybern. Part C **34**(2), 210–218 (2004)