IEEE/RSJ International Workshop on Intelligent
Robots and Systems IROS '91, Nov. 3-5, 1991,
Osaka, Japan. IEEE Cat. No. 91TH0375-6

Perception of Rugged Terrain for a Walking Robot: True Confessions and New Directions

Regis Hoffman and Eric Krotkov
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA. 15213 U.S.A.

Abstract

The Planetary Rover Project at Carnegie Mellon University has developed and demonstrated the Ambler, an autonomous, legged mobile robot that operates in rugged environments. During the course of the Ambler walking experiments unexpected shortcomings in the robot's perception system were noted. Unusual perception sensor behavior, random noise sources, and interaction of the perception system with other system components caused both minor errors and large failures in the system. This paper analyzes the sources of the perception system errors, presents solutions to these problems, and proposes a design for a more robust perception system.

Introduction

The Ambler is an autonomous, orthogonal walking robot that operates in rugged environments. The mechanism has two stacks and six circulating legs [2] [1] . On top of the Ambler is mounted a laser scanner to obtain 3-D data of the local terrain (Figure 1).

Figure 1: The Ambler Robot

The three major components of the Ambler's software system are perception, planning, and control. The perception system builds elevation maps of the terrain from 3-D range images. The planning system uses elevation maps to determine obstacle-free routes and to select good footfall locations. The control system executes body and leg trajectories generated by the planner.

During the course of walking experiments, the perception system built maps that permitted the Ambler to negotiate extreme terrain. More than 1000 m of terrain was traversed, accumulated during the course of 100 trials. With this success came the discovery of unexpected shortcomings of the perception system due to inaccuracies in the computed elevation map. The errors encountered were both spatial (large elevation spikes and incorrect elevation values within a small region), and temporal (divergence of meter-sized terrain patches from the ground truth over time). Spatial map errors (such as elevation spikes) introduce large obstacles that the planner considers as blocking the robot's path - in the worst case, the planner finds all paths blocked. Temporal errors lead to sometimes dramatic collisions of the robot with the terrain, as the perception maps have diverged significantly from ground truth.

A number of research efforts have addressed the problem of mapping rugged, outdoor terrain for mobile vehicles. Researchers have investigated sensors ranging from passive vision to sonar to structured light to laser scanners, and various combinations. These efforts have led to implementations on real robots, including (in alphabetical order) Attila [3] , Hilare [4] , Navlab [7] , Robby [8] , and many others. A common conclusion about the results of many of these efforts is that fielded perception systems often do not work as well as expected (despite what the advertising may say).

Our immediate goal for the Ambler project is long-duration (> 5 hour) autonomous walking - this demands reliable and robust perception systems. Our current achievement is about 45 minutes of autonomous walking. This paper appraises the implemented perception system, emphasizing situations where the perception system failed (true confessions). We then present solutions implemented to solve several of these problems (new directions), and finally argue that the open-loop nature of the perception system must be replaced by a closed-loop system, where feedback from leg-terrain contact is added to the system to enable long-duration autonomous walking.

Perception System

The perception system builds and maintains terrain elevation maps of the 3-D environment [5] . The elevation maps are computed from range images acquired by a Perceptron laser scanner. The scanner provides range and reflectance images of size 256 x 256 pixels, with 12 bits of data per pixel, at a frame rate of 0.5Hz. The scanner has a 60 degree field of view in the horizontal and vertical directions, and an operating range of 2-40 meters [6] . Figure 2 shows a pair of range and reflectance images from the Ambler test course. The scene consists of a sand base with 1-2 meter high boulders (foreground) and a wooden ramp (background).

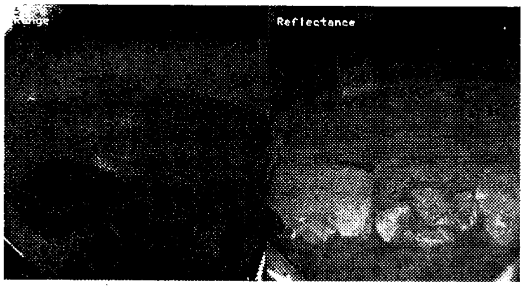

Figure 2: Range (left) and Reflectance (right) Image

In the simplest case, the perception system takes a range image and computes an elevation map of the local terrain. As the robot moves, a sequence of range images I_0, I_1, ..., I_N together with the position of the robot body B_0, B_1, ..., B_N (in a global reference frame) is taken. Given this set of range images and the position of the robot when the images were taken, we compute a map of the environment by merging maps created from each single image. If a point in the map cannot be computed (due to invalid range data, outside the field of view, or occlusion), the map point is tagged as *unknown*.

Figure 3 shows a 4 x 5 meter elevation map at 10 cm resolution computed from the range image in Figure 2. The dark areas are labeled unknown regions, either because of terrain occlusion or bad data in the range image.

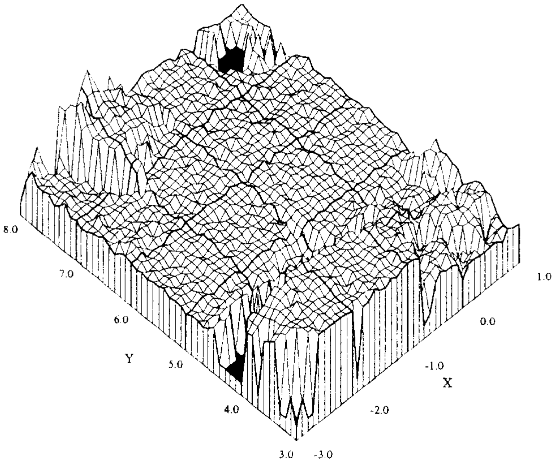

Figure 3: Elevation Map

Implicit in the design of the perception system were the following assumptions:

- The quality of a single range image was high. Specifically, all pixels represented valid range measurements (with some noise), and were relatively independent of surface material.
- A sequence of range images of a static scene would vary little over time, and would not be seriously affected by environmental changes such as temperature and ambient light.
- The laser scanner's position in the world frame was accurately known.

This paper discusses why these assumptions were invalid, and how this affected the performance of the perception system. The duality between images and maps implies that errors in the image result in errors in the map; however, not all map errors are caused by image errors.

Phantom Rocks

The elevation map shown in Figure 4 appears reasonable, with some relatively flat terrain nestled next to a rock of height 50-60 cm. However, the scene consists solely of flat terrain. The reason for this phantom rock is an internal reflection from the exit window of the laser scanner. The left-hand side of Figure 5 shows the reflectance image of the scene that yielded the map in Figure 4, with the outlined region indicating the map bounds. After close inspection of the corresponding range image, it was determined that the range measurements were *not* constant in this region, but varied quadratically over a 40 x 20 pixel region, and appear as a smudge in the range image. The right-hand

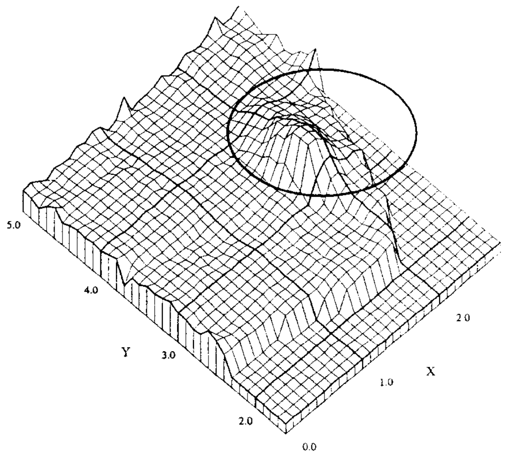

Figure 4: Phantom Rock (circled)

side of Figure 5 shows a magnified portion of the lower right corner of the range image - the smudge appears as a darker region. Because all the range values in this smudge area are less than the true value, they appear closer to the scanner, and become small hills or rocks in the elevation map.

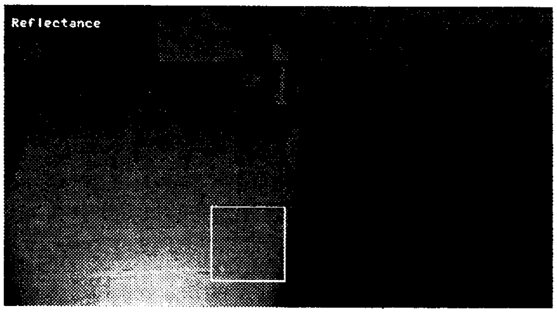

Figure 5: Reflectance Image (left) and Corresponding Range Image (right, magnified).

Because the smudge is always at a fixed location in the image, we threshold the range image in the region where the smudge occurs, and reject pixels that are not within the threshold. These pixels are labeled invalid, and are not used in the computation of the elevation map. The disadvantage to this approach is that more points in the elevation map are marked as unknown, as less range data is available. The map in Figure 6 shows the result of this masking. Most of the map values corresponding to the incorrect image pixels are labeled unknown, and must be recomputed from the next image taken.

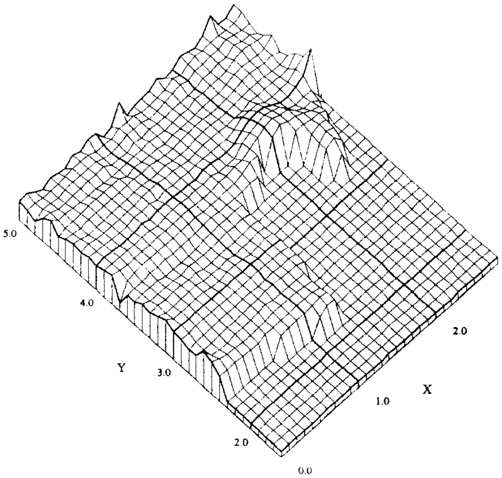

Figure 6: Phantom Rock Reduced by Masking.

Map Spikes

Map spikes are large, relatively isolated values in the elevation map. Map spikes may result in planning system errors, as positive spikes appear as impassible obstacles, and negative spikes are craters or holes to be avoided.

The map in Figure 7 has a region of large negative elevation values (approximately -20 meters) whereas the terrain actually is 1-2 meters high. It is possible to remove most spikes by a median filter, but elevation spikes sometimes occur in clusters, in which case they may not be eliminated by the median filter. Also, it is important to understand the causes of map spikes, as they may indicate a fundamental problem with the perception system.

The map spikes generally occur at the border of images - this led to the discovery of yet another problem with images. As seen in both Figure 2 and Figure 5, the lower corners of the images exhibit a black region (or vignette) that contain incorrect range values. This vignette is caused by the exit window of the laser scanner, which occludes a small portion of the field of view. In this region, range values are again incorrect and must not be used in the elevation map computation. As with the image smudge, this vignette occurs at a fixed location in the image, and we apply a threshold operation to mark the offending pixels as invalid. The map after the vignette pixels are identified and excluded from the map computation is shown in Figure 8.The major difference between the two maps is that regions that were previously marked as having valid but very large values are now marked as unknown.

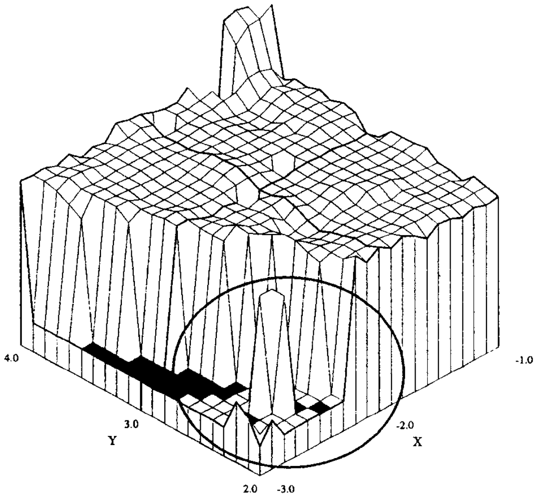

Figure 7: Negative Elevation Spikes (circled)

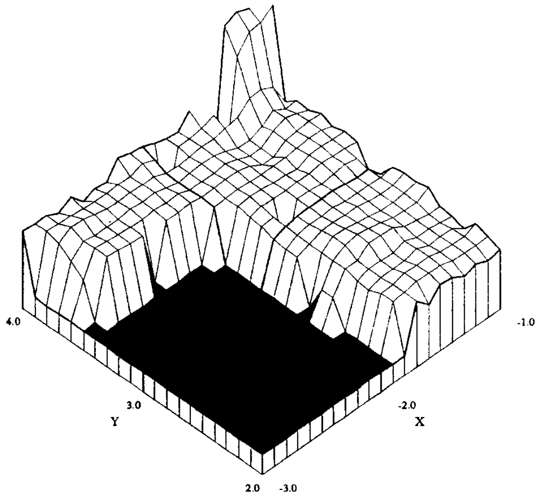

Figure 8: Spikes Reduced by Image Mask.

Material Properties

In theory, the range value returned by the sensor is largely independent of the surface material. In practice, this is rarely the case. Two materials with different reflectance properties placed an equal distance from the range sensor may yield different range measurements. This has a severe impact on elevation maps constructed across terrain consisting of varying materials. This effect is illustrated in Figure 9.

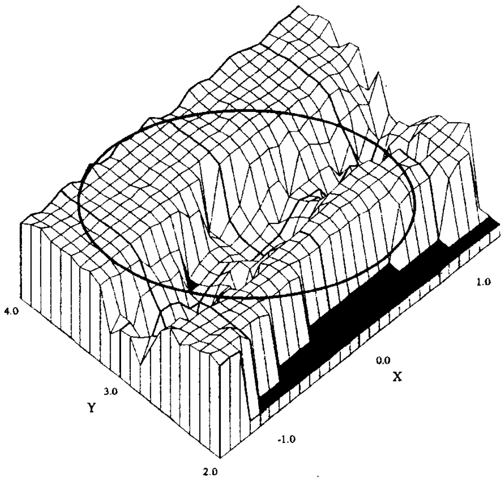

Figure 9: Map Discontinuity Across Surface Boundary (circled).

The trough in the map (about 70 cm lower in height than the background terrain) is purely an artifact of a change in surface materials, *not* a product of the terrain geometry. The area to either side of the trough is a light color material (concrete), and the trough is a black, grease covered band running across the concrete. The different surface material properties yield different range values, and hence produce variations in the elevation map.

There is no general solution to this problem. It is possible to perform a connected region analysis on the range image and reject all pixels in the image that are not smoothly connected (since dark materials we have studied have range values 20-30 cm less than lighter materials, they will not be smoothly connected to adjacent points in the scene that consist of lighter color materials). This can sometimes lead to the removal of an unacceptable number of pixels from the image, and leads to a large number of points in the map being classified as unknown.

Figure 10 shows the result of applying the connected region analysis to identify invalid pixels in the range image (pixels that arise from dark surface materials). This operation was applied to the same image that was used to compute the map from Figure 9, and shows that it is possible to minimize the effects of the surface material problem at the expense of labeling a possibly significant number of map points as unknown. Note the circled region in Figure 10 - there are now two unknown points where before all the elevations were known (in addition to the large unknown band at $y < 3$). This causes the planning system to act more conservatively in this region.

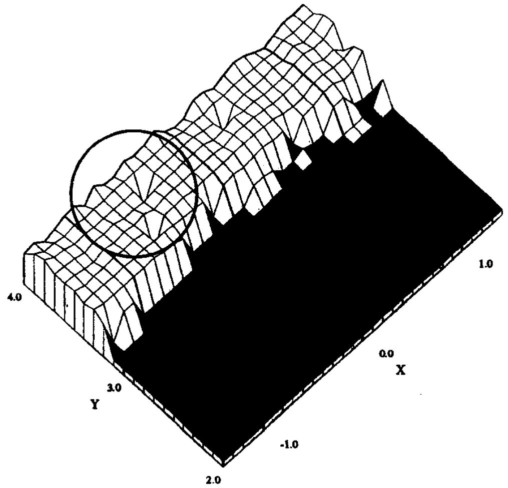

Figure 10: Area from Figure 9 Using Connected Region Analysis to Remove Invalid Pixels

Temperature Drift

The initial version of the laser scanner was very sensitive to ambient temperature variations. Changes in temperature from 65 degrees F to 85 degrees F produced changes in the range measurement of up to 3 meters. The large range drift lead to calibration errors, and the computed maps diverged drastically from ground truth. This was the principal cause of perception system failures during the walking experiments.

It is difficult to compensate for such large fluctuations in temperature. An upgrade to the sensor electronics and cooling system has significantly reduced this temperature dependence. Figure 11 shows the current range drift over time. A range measurement to a fixed target was acquired over 800 trials during a two hour period. The sensor was subject to an ambient temperature range of 80 degrees F (at Trial 0) to 100 degrees F (at Trial 180), and returned to 82 degrees F (at Trial 250). The graph records the range measurement versus time (or trial). The sensor has about 15-20 cm noise with each measurement, and up to 20 cm range drift over this temperature. The range measurement still varies, but not nearly as much as in the original version of the sensor. We expect that this sensor modification will permit longer duration walking.

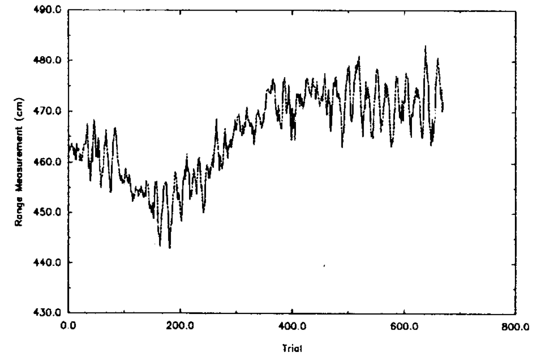

Figure 11: Range Drift over Time

Map Seams

An elevation discontinuity or seam of 10-15 cm is illustrated in Figure 12.

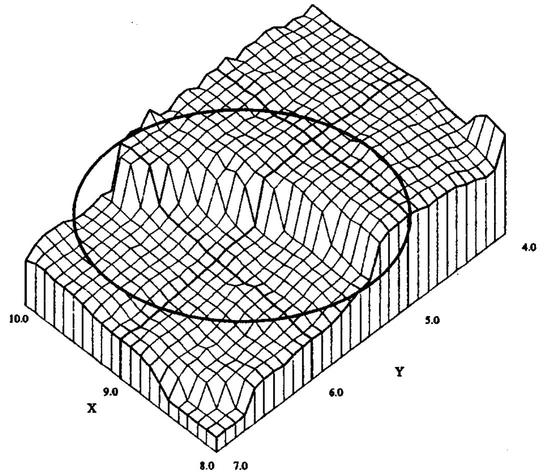

Figure 12: Map Seam (circled).

The seam did not arise from an error in the range image, but is a result of a dead-reckoning position error. The map in Figure 12 was constructed from two range images. The map region with the lower elevation values was computed from one image, and the higher elevation region was computed from the second range image. Recall that a map is built from the range image I and the body position associated with the image B_i. Any error in the body position will result in errors in the map. The seam in the map is caused by a change in the body height (according to dead-reckoning) between the two images, even though the body height actually remained constant. The height of the seam (about 10-15 cm) was the error in the two dead-reckoned body positions.

Discussion

The longest duration walking experiment lasted 45 minutes (30 steps) before inaccuracies in the perception system elevation maps led to either robot-terrain collisions or decisions by the planning system to halt further advance due to impassable terrain obstacles. Map spikes, phantom rocks, and divergence of the computed maps from ground truth all caused significant problems.

Initially, errors in the range images and elevation maps were detected and corrected by the experimenters. However, completely autonomous systems must identify and correct problems independent of human assistance. Can a map seam be distinguished from a true elevation discontinuity? Can phantom rocks be detected? Can we compensate for sensor calibration errors due to a temperature drift? Our major conclusion from this work is that perception cannot operate open-loop - there must be some feedback from the robot's physical contact with the environment. Without feedback it is difficult to assess the quality of the maps generated by the perception system. For example, as the robot walks through the terrain, the height of the legs when they make terrain contact (the ground truth) can be compared to the elevation computed by the perception system. Any divergence between the two indicates an error in the perception system that possibly can be corrected.

From the experience with the perception system, a new system was designed and implemented. It incorporates all the techniques discussed above, and monitors ground truth measured by the robot's legs and compares it to the perception system maps to detect errors. Walking experiments are now being conducted to compare the two systems.

Acknowledgments

The authors would like to thank Martial Hebert who made important contributions to the design and implementation of the initial perception system, Henning Pangels and Kenichi Arakawa for assistance in calibration, and Mike Blackwell for work on the Perceptron scanner.

Bibliography

[1] J. Bares, M. Hebert, T. Kanade, E. Krotkov, T. Mitchell, R. Simmons and W. Whittaker. Ambler: An Autonomous Rover for Planetary Explorations. *IEEE Computer.* 18-26, June, 1989.

[2] J. Bares. *Configuration of Autonomous Walkers for Extreme Terrain.* Ph.D. thesis, Department of Civil Engineering, Carnegie Mellon University, 1991.

[3] R. Brooks, P. Maes, M. Mataric, and G. More. Lunar Base Construction Robots. *Proceedings IEEE International Workshop on Intelligent Robots and Systems*, pages 389-392. Tsuchiura, Japan, July, 1990.

[4] G. Giralt, R. Chatila and M. Vaisset. An Integrated Navigation and Motion Control System for Autonomous Multisensory Mobile Robots. Brady, M. and Paul, R. (editor), *Proceedings First International Symposium on Robotics Research.* MIT Press, Cambridge, Massachusetts, 1987.

[5] M. Hebert, T. Kanade, E. Krotkov and I.S. Kweon. Terrain Mapping for a Roving Planetary Explorer. *Proceedings IEEE Robotics and Automation*, pages 997-1002. Scottsdale, Arizona, May, 1989.

[6] I. S. Kweon, R. Hoffman and E. Krotkov. *Experimental Characterization of the Perceptron Laser Rangefinder.* Robotics Institute Technical Report CMU-RI-TR-91-1, Carnegie Mellon University, 1991.

[7] C. Thorpe (editor).*Vision and Navigation: the Carnegie Mellon Navlab.* Kluwer, 1990.

[8] B. Wilcox and D. Gennery. A Mars Rover for the 1990's. *Journal of the British Planetary Society.* 40 484-488, 1987.