Surface Roughness Measurement for Outdoor Mobile Robotic Applications

Alan M. Zhang and R. Andrew Russell
Centre for Perceptive and Intelligent Machines in Complex Environments:
Intelligent Robotics
Monash University, Clayton
VIC 3800, AUSTRALIA
{alan.zhang, andy.russell}@eng.monash.edu.au

Abstract

The ability to measure terrain roughness is essential for safe and efficient navigation of outdoor mobile robots. This paper describes a prototype system consisting of a colour camera and a laser stripe generator for the measurement of ground surface roughness. With outdoor mobile robotic applications in mind, the system provides a cheap, fast and robust alternative to other methods of roughness measurement. In this paper, details are given of the hardware configuration, the method of roughness characterisation, a simple navigation strategy demonstrating the ability of terrain roughness measurements to aid in navigation, and some preliminary results.

1 Introduction

Autonomous mobile robotic navigation systems operating in natural, unstructured environments require an assessment of the traversability of the surrounding terrain and this could be based on surface roughness. One approach to this problem could involve classification and pattern recognition in which passive vision is the main sensory input. Regions in a scene are classified as either smooth or rough depending on image features such as colour and visual texture. Outdoor scenes are often subject to large variations in illumination and consequently large variations in colour of the same surface [Austin and Barnes, 2003]. Hence, these classifier systems often require supervised training on large volumes of sample data under different illumination conditions and surface types. Their behaviour for novel surface types can be unpredictable. Other systems measure surface roughness directly without the need for inference using prior knowledge. The stereo vision system of McDonald, *et al.* [1999] uses disparity information to measure surface roughness. Systems based on disparity are affected by surface texture or the lack of texture and tends to have a low resolution in detecting surface undulation. Radar has been used extensively for numerous geologically related remote sensing applications. However, the radio wavelength limits the resolution of such systems. Surface material properties, such as the dielectric constant, also have a large effect on the characteristics of the radar backscatter [Inggs and Lord, 2000]. Road roughness laser profilometers use a spot laser and a camera to triangulate the elevation of the road surface [Sayers and Karamihas, 1996]. On-board accelerometers measure the vertical displacement of the system and wheel odometry measures the horizontal displacement. A cross-sectional elevation profiles of the surface is produced by integrating the elevation found by triangulation with the system displacement measurements. This method not only requires very accurate vertical acceleration and odometry measurements, it is also slow in constructing a surface profile. Lidar based systems are generally bulky and expensive.

This paper describes a prototype system for ground surface roughness estimation intended for outdoor mobile robotic applications. Using a camera/laser stripe combination the system provides near instantaneous cross-sectional surface profile by triangulation. Although much research has been conducted in the area of roughness measurement using a camera/laser stripe combination, such as in [Darboux and Huang, 2003; Kim and Chang, 1999; B. Xu and D. F. Cuminato, 1998; Norbert H. Maerz, et al., 2001] and many more, these systems were not designed to be used on mobile robots. Our system was designed to over some of the challenges and requirements poses by outdoor mobile robotics. It is designed to be robust against different illumination conditions, surface colour textures and mechanical jitter, is low in weight and inexpensive. Organisation of this paper is as follows: Section 2 presents the assumptions made and the system architecture. Section 3 outlines the process of extracting the elevation profile from images and the method of estimating roughness. Section 4 describes the sub system that utilises the rough-

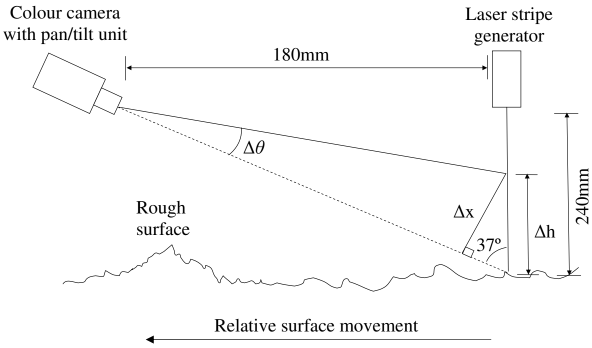

Figure 1: Hardware configuration consists of a colour camera mounted on a pan/tilt unit (not shown) and a laser stripe generator.

ness estimate to generate traversability maps. Section 5 considers some future directions of this research.

2 System Architecture

2.1 Assumptions

We assume that the measured surfaces are statistically isotropic, such that vertical cross-sections of the surface in any direction would have identical statistical properties. This assumption simplifies the problem of characterising surface roughness into a two dimensional problem from a three dimensional one. It is expected that this assumption will hold for most natural surfaces. It is also assumed that the surface material does not reflect light in a highly specular fashion. Some surfaces could have high specularity when wet, but most surfaces of interest have good diffuse reflection properties.

2.2 Equipment Setup

The system configuration is shown in Figure 1. A 1mW, 650nm red laser was spread into a stripe using a cylindrical lens. The inexpensive pinhole type colour camera used in the system had a resolution of 320 by 240 pixels. The focal length was found to be 450 pixels through camera calibration. Images were captured at 10 frames per second. Lens distortion was minimal so no compensation was required for the captured images. The camera was mounted on a pan/tilt unit so that it could obtain a view of the surroundings as well as looking directly at the laser stripe. A computer controlled the pan/tilt unit and the camera via a USB interface. Camera shutter speed and gain were adjusted automatically by the computer to accommodate the change in lighting conditions in outdoor environments. The stripe generator and camera assembly were fixed onto a rigid structure to maintain the geometry shown in Figure 1. The entire system was mounted at the front end of a mobile robotic platform such that the stripe was perpendicular to the robot's direction of motion. It was positioned so that when the robot moved, the surface directly in front of the robot was scanned. Our system, unlike most road profilometers, did not require odometry measurements or sensors to measure vertical displacement.

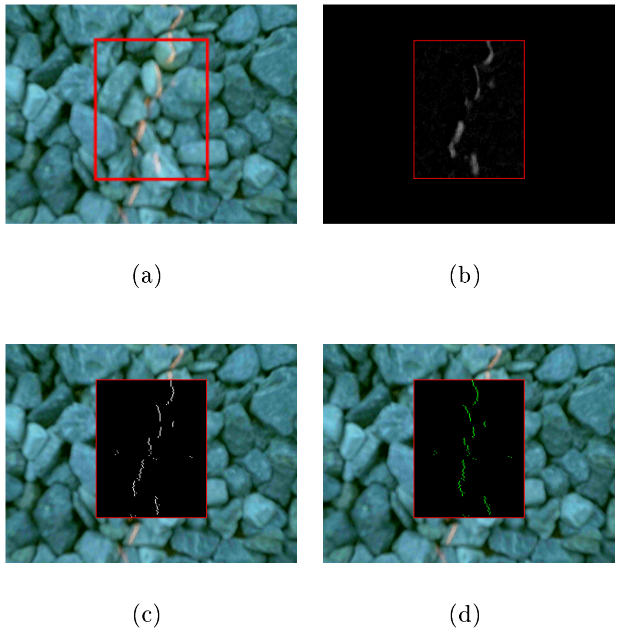

Figure 2: Stages in the extraction of elevation profile. (a) Original captured frame, the rectangle in the middle of the image is the region where the stripe is expected. (b) Red channel after frame differencing. (c) Extracted elevation profile. (d) Elevation profile with trend removed.

3 Roughness Measurement

3.1 Extraction of Elevation Profile

To measure surface roughness, the camera/laser assembly was moved over the surface to obtain a scan of the surface. Due to its very low power, the laser stripe in the camera image needed to be enhanced. Narrow band colour filters are good at enhancing the visibility of the laser stripe [Darboux and Huang, 2003], but surface colour information would also be lost. Therefore a two frame difference was performed to enhance the laser stripe while the system is scanning across the surface instead of a colour filter. In the differencing process, the movement of the surface between the current and previous frame had to be measured. We took a patch of surface texture from the current frame and found the corresponding patch in the previous frame using sum of magnitude difference as the distance measure. The resulting disparity was taken as the movement of the surface between frames. The previous frame was shifted by an amount equal to the measured disparity so that the surfaces in the two images were lined up. A pixel-wise

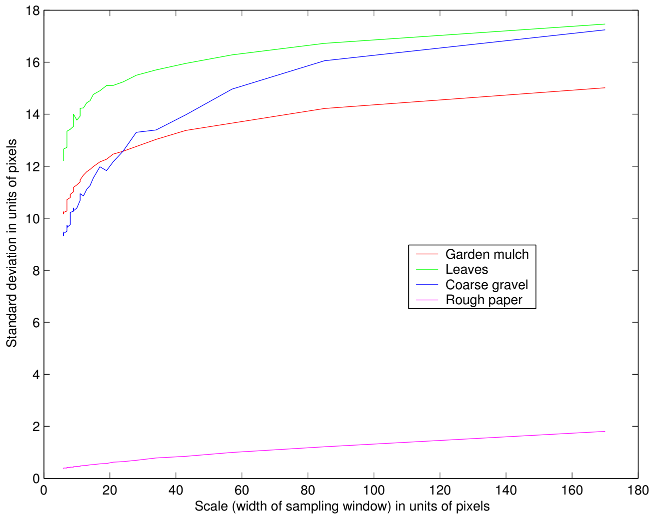

Figure 3: Multiscale variance of different surfaces averaged over 249 frames. Variances of each scale corresponding to different sampling window positions were averaged to provide a variance estimate for that scale

subtraction was then performed on the images. The result is the enhancement of the laser stripe and suppression of surface colour texture. Figure 2(a) shows the original image and Figure 2(b) shows the result of frame differencing in the red channel. Note that frame differencing was only performed inside a window where the laser stripe was expected. The position of the stripe was then extracted by taking the *maximum* intensity in the red channel along each horizontal scan line. No thresholding was used because surfaces might be composed of different materials that reflect light to varying degrees. Thresholding would erroneously reject genuine but weak laser reflections.

Assuming that $\Delta\theta$ in Figure 1 is small, the displacement of the laser stripe is directly proportional to the vertical undulation Δh. The extracted stripe is therefore a scaled version of the cross-sectional elevation profile of the surface. With a camera focal length of 450 pixels, the system was capable of resolving surface undulation down to approximately 1.2mm under ideal operating conditions. Figure 2(c) shows the extracted elevation profile.

3.2 Measure of Roughness

Statistical measures of roughness have been studied extensively in the literature [Rouillard and Sek, 2002; Arakawa and Krotkov, 1993], especially in the area of geological remote sensing [Manninen, 1997]. A common measure is the standard deviation (known as RMS height) and autocorrelation length used together as a pair to characterise surface profile. These measures assume that the surface characteristics are stationary. Naturally occurring surfaces are often non-stationary, in that the standard deviation and autocorrelation length vary as the resolution of the profile is changed. A multi-scale method of characterisation is therefore needed [Manninen, 1997]. Many natural surfaces have been shown to exhibit fractal characteristics, but only over a limited range of spatial scales [Arakawa and Krotkov, 1993]. Alternatively, the power spectrum would completely describe the characteristics of a surface at all spatial scales but it requires the processing-intensive Fourier transform. It is also prone to aliasing and sensitive to outliers. Our system used a multiscale extension of the RMS method and is described below.

A sampling window of a fixed width was passed over the elevation profile to obtain sets of samples. The variance of each set was calculated and associated with the position of the sampling window. The sampling window was then changed to different widths and variances were calculated for each window width and position. The *spatial scale* refers to the width of the sampling window for which the variances were calculated. So for each *scale*, we have a variance measure for each *position* of the sampling window. Vertical displacement and roll of the system relative to the surface often add a constant slope and offset to the measured elevation. This effect is particularly pronounced when a small sized robot without wheel suspension travels over rough terrain, as is the case with the robotic platform used in this research. Mechanical vibrations also contributed to the distortion of the laser stripe. Least squares linear regression was used to remove trend from samples at *each* scale before the multiscale variance calculations. Figure 2(d) shows the elevation profile after the trend is removed from the largest scale (i.e. the entire width of the profile). The final result is the characterisation of the elevation profile by a set of variances, which indicate the roughness at different spatial scales.

The system was tested on surfaces that outdoor mobile robots are likely to encounter such as grass, garden mulch and gravel. It was discovered that these surfaces often caused severe occlusions of the laser stripe. i.e. the camera could not see parts of the stripe because it was blocked by surface protrusions. Therefore, no measurements could be made on the elevation of the occluded parts of the stripe. Recall that the elevation profile extraction process described in Section 3.1 simply finds the maximum intensity reading in the red channel along each horizontal scanline and that frame differencing removes surface colour texture, the gaps would then be effectively filled with random measurements. This is a nice side-effect of not using thresholding because occlusion is an indication that the surface is rough. With the roughness measurement method described above, a smooth surface with occlusion is likely to be measured as rougher, due to the random measurements, than a

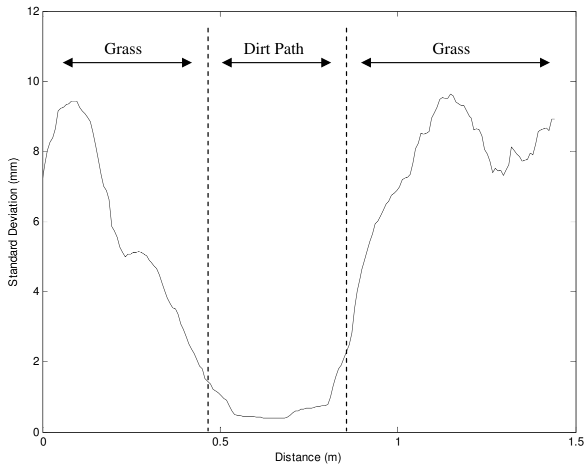

Figure 4: Standard deviation of surface elevation at a scale of 100mm taken while crossing a dirt path in a grass area. The dirt path can be seen to extend from 0.5m to 0.8m.

smooth surface without occlusion.

A running average of individual roughness measures of surface profiles was used to reduce the effect of spurious measurements. Also note that only 2 frames were needed to provide an elevation profile, so the system could be said to provide a near *instantaneous* measurement of surface roughness (though this measurement tends to be noisy).

3.3 Results

Results of measurements taken from a few different types of surfaces, averaged over 240 frames for each surface, and over a range of scales are shown in Figure 3. Variances for a single scale corresponding to different sampling window positions were averaged to provide a variance estimate for that scale. Because of the way the system handled occlusions, the roughness measurements were affected by both the amount of occlusion and the roughness of the visible parts of the surface. Thus, the measurements could be interpreted as *beliefs* of surface roughness. Nevertheless, Figure 3 clearly indicates that the multiscale method was able to distinguish between surfaces of different roughness. The surface referred to as *rough paper* in Figure 3 was made by scrunching a piece of paper and then flattening it out. Notice how the roughness of coarse gravel drops quickly as the scale decreases. This is due to the fact that surfaces of individual rocks are rather smooth.

Figure 4 shows the roughness measurements taken when the robot crossed a dirt path in the middle of a grass lawn (dirt path shown in Figure 5). Result shown in Figure 4 is a 20 frame running average and at a scale of 100mm. The low roughness from a distance of 0.5m to 0.8m is the dirt path.

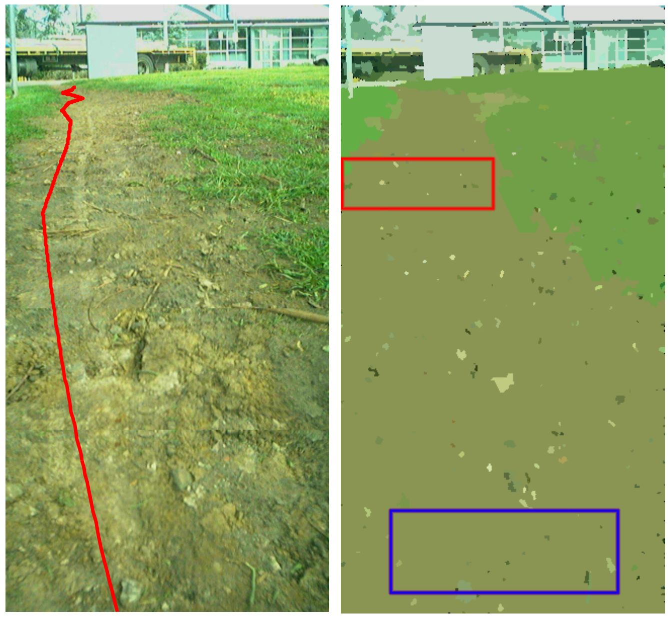

(a) (b)

Figure 5: (a) Mosaic image used to create the traversability map. The red line is the trajectory the robot subsequently followed. (b) Segmented image. Regions are displayed with their average colour. The region with the maximum number of pixels inside the rectangle near the bottom of the image was associated with the roughness just measured. The rectangle near the top is a patch of ground approximately 1.5 meters away identified as traversable.

4 Ground Traversability Map

This section describes a system that utilises the roughness measurements obtained in the last section to infer traversability of the surrounding terrain

4.1 Segmentation

Once the roughness of a surface has been measured the camera's elevation angle was changed in steps of 15 degrees using the pan/tilt unit, starting from looking directly at the surface just measured (with the laser turned off), until the camera is oriented horizontally. An image was captured for each angle. These images were then combined to produce a mosaic such as the one shown in Figure 5(a). The bottom of Figure 5(a) is the area whose roughness has just been measured. To reduce the effect of noise, the SNN edge preserving filter [Chen and Shih, 2002] was applied over the entire mosaic. Segmentation of the smoothed mosaic was performed using the Colour Structure Code algorithm (CSC) described in Rehrmann and Priese [1998]. The CSC algorithm is essentially a merge-and-split method where each pixel was initially treated as a separate region. Regions were then merged based on their colour similarity. When no more merging was possible, a split stage broke up regions with smoothly varying colours that had too much colour variance in the same region. The colour similarity predicate function of the CSC algorithm operates in the HSV colour space. It employed a lookup table to determine colour similarity thresholds to cater for large variances in illumination encountered in outdoor scenes. For details of the algorithm and its implementation refer to Rehrmann and Priese [1998]. Figure 5(b) shows the segmented image in which regions are displayed with their average colour. The region with the maximum number of pixels inside the rectangle near the bottom of Figure 5(b) was assumed to have the same roughness as the surface just measured.

4.2 Navigation

A simple behaviour that used the traversability map was implemented on a mobile robot to demonstrate the viability of this approach. The algorithm uses the flat earth assumption and looks for a patch of ground approximately 1.5 meters away that was classified as smooth and wide enough for the robot to pass through. The robot would then adjust its heading towards this patch of ground and move forward approximately 1.5 meters. If no suitable patch was found, the robot would rotate in the most promising direction until it finds such a patch. In Figure 5(b), a suitable patch was found as indicated by the rectangle near the top of the image. This simple system was able to follow the dirt path shown in Figure 5(a) for more than 10 meters before the dirt path ended. The trajectory of the robot recorded by odometry is superimposed on Figure 5(a) in red. The navigation system and the roughness measurement system are currently separate, but when integrated, it would allow an autonomous mobile robot to explore its environment safely and efficiently by traveling on relatively smooth surfaces.

5 Future work

The major limitation of the system is the laser stripe generator. The 1mW laser was too weak to be visible when the surface to be measured was under strong lighting, such as direct sun light. A more powerful laser or combining a number of laser stripes would enable the system to work under a greater variety of lighting conditions. The least squares linear regression method used to remove the trend in the elevation profile was sensitive to outliers. Currently, a running average of measurements over multiple frames was used to reduce the effect of these outliers. A robust linear regression method would yield more stable measurements, hence reduce the number of readings needed in the running average and consequently improve system response. The ability to recognise different types of surfaces is beneficial because certain surfaces are pliable, such as grass, and do not hinder the robot too much even though the measured roughness is high. Multiscale variances of various surfaces could be compiled into a template database to allow classification of surfaces. An advantage of not using a colour filter to enhance the laser stripe is that rich colour information can be collected at the same time as roughness measurements are made, providing an opportunity to combine colour and visual texture with the roughness measurements to obtain more reliable surface classification. The system has only been tested on a limited number of surfaces of relatively low roughness, but results thus far indicate that it is robust against mechanical vibration and jitter of the mobile platform. Jitter and vibration leads to motion blur in the captured images. For surfaces with low roughness, this has not been observed to be a big problem, but further testing on much rougher surfaces needs to be conducted to verify this claim.

6 Conclusion

We have demonstrated a prototype system for surface roughness measurement consisting of only a laser stripe generator and a colour camera. The system was designed for outdoor mobile robotic applications. It is low weight, low cost, provides near instantaneous cross-sectional surface elevation profile, continuously takes measurements while the robot is moving, provides roughness measures at different spatial scales, and is insensitive to mechanical jitters and vibration due to the unevenness of the ground. A system that generates ground traversability maps was developed to demonstrated the potential use of the roughness measurement system.

Acknowledgments

The work described in this paper was supported by the Australian Research Council funded Centre for Perceptive and Intelligent Machines in Complex Environments. We would like to thank Geoff Taylor for his help in setting up the equipment and valuable discussions; Lutz Priese, Volker Rehrman and Patrick Sturm for making the source code of the CSC algorithm freely available online.

References

[Arakawa and Krotkov, 1993] Kenichi Arakawa, Eric Krotkov. Estimating Fractal Dimension of Natural Terrain from Irregularly Spaced Data. *Proceedings* *of the 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems,* July 26-30, 1993

[Manninen, 1997] A. T. Manninen. Multiscale Surface Roughness and Backscattering. *Progress in Electromagnetics Research, PIER 16,* 175-203, 1997

[McDonald, et al., 1999] A. J. McDonald, S. Crossley, J. C. Bennett, S. C. M. Brown, G. Gookmartin, K. Morrison, S. Quegan. Stereo vision measurements of soil surface characteristics and their use in model validation. *Committee on Earth Observation Satellites 1999 SAR workshop.* 1999.

[Darboux and Huang, 2003] Frederic Darboux and Chihua Huang. An Instantaneous-Profile Laser Scanner to Measure Soil Surface Microtopography. *Soil Sci. Soc. Am. J.* 67:92-99 (2003).

[Rouillard and Sek, 2002] Vincent Rouillard and Michael A. Sek. A statistical model for longitudinal road topography. *Road & Transport Research, Vol 11, No 3,* September 2002.

[Howard, et al., 2001] Ayanna Howard, Edward Tunstel, Dean Edwards, Alan Carlson. Enhancing Fuzzy Robot Navigation Systems by Mimicking Human Visual Perception of Natural Terrain Traversability. *Joint 9th IFSA World Congress and 20th NAFIPS International Conference, Vancouver, B.C., Canada,* July 2001, pp. 7-12

[Rehrmann and Priese, 1998] Volker Rehrmann and Lutz Priese. Fast and Robust Segmentation of Natural Color Scenes. *Proceedings of 3rd Asian Conference on Computer Vision (ACCV'98) R.Chin and T.C.Pong (Eds.).* 1998.

[Chen and Shih, 2002] Shawn Chen and Tian-Yuan Shih. On the evaluation of edge preserving smoothing filter. *Proceedings of Geoinformatics (CD), June 1-3, Nanjing, China, paper C43.* 2002

[Inggs and Lord, 2000] M. R. Inggs and R. T. Lord. Applications of Satellite Imaging Radar. *South African Institute of Electrical Engineers, SAIEE, South Africa,* 2000

[Sayers and Karamihas, 1996] M.W. Sayers and S.M. Karamihas. Interpretation of Road Roughness Profile Data. *University of Michigan Transportation Research Institute Report UMTRI-96-101,* 1996

[Austin and Barnes, 2003] David Austin and Nick Barnes. Red is the new black - Or is it? *Proceedings of the 2003 Australasian Conference on Robotics & Automation, Brisbane, Australia,* December 1 - 3, 2003

[Kim and Chang, 1999] Dea-Gyu KIM and Won-Seok CHANG. A Study on a 3-D Profilemeter Using Dynamic Shape Reconstruction with Adaptive Pattern Clustering of the Line-Shaped Laser Light. *TENCON IEEE, 1999*

[B. Xu and D. F. Cuminato, 1998] B. Xu and D. F. Cuminato. Evaluation of Fabric Smoothness Appearance Using A Laser Profilometer. *Textile Research Journal, 68(12), 900-906, 1998*

[Norbert H. Maerz, et al., 2001] Norbert H. Maerz, Antonio Nanni, John J. Myers and Greg Galecki. Laser Profilometry for Concrete Substrate Characterization Prior to FRP Laminate Application. *Concrete Repair Bulletin May/June 2001*