

Laser Intensity-Based Obstacle Detection

John Hancock, Martial Hebert, and Chuck Thorpe
[jhancock, hebert, cet]@ri.cmu.edu

Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract

We present a novel method for obstacle detection for automated highway environments. Laser range scanners have frequently been used for obstacle detection for mobile robots. Although most laser scanners provide intensity information in addition to range, laser intensity has been ignored by most researchers. We show that laser intensity, on its own, is sufficient (and better) for detecting obstacles at long ranges in mild terrain such as an automated highway.

1 Introduction

Highway obstacle detection is a challenging problem. Highways present an unknown and dynamic environment with real-time constraints. In addition, the high speeds of travel force a system to detect objects at long ranges. While a variety of competing methods have been proposed for on-road obstacle detection most of the work has focused on detecting large objects, especially other vehicles. There are a number of methods that can successfully detect moving vehicles but the more difficult problem of finding small, static road debris such as tires or crates remains unsolved. Systems such as AHS which demand high levels of safety are not feasible unless these critical problems are addressed[1]. We propose to use laser reflectance to detect small obstacles at high vehicle speeds in automated highway applications.

Laser range scanners have been used for obstacle detection for a number of years, especially for navigation in rough terrain. Laser range scanners operate by sweeping a laser over a region of interest and measuring, at each pixel, the time it takes for the laser to leave and return to the sensor. Since the speed of light is known, the distance to every pixel can be calculated. Most laser range scanners also provide the intensity of the returned signal at every pixel. However, this second piece of information, often referred to as the reflectance, has essentially been ignored by most researchers.

Traditional methods of laser processing for off-road obstacle detection usually involve 3 steps. First, using the known pose of the sensor, range points are transformed into elevations in a map. In image-based techniques, the elevation map is generally sparse and noisy, and smoothing is generally necessary before a terrain-typing algorithm is run to classify areas of the elevation map as navigable or unnavigable. In more recent pixel-based techniques, the map is updated pixel-by-pixel, and the local map containing discrete obstacles is scrolled with vehicle motion to make better and faster (lower latency) use of new and old data. Both processes are complicated and these methods have another drawback in that elevation maps may not represent obstacles well. Discrete grid size and cell elevation averaging make it impossible to represent vertical planes[5],[7]. Range measurement devices have additional problems such as inaccuracies and significant noise at low signal levels, and range ambiguities.

On the other hand, laser reflectance ought to provide us with a more direct means of finding obstacles or vertical surfaces. At the long lookahead distances and grazing angles typical of high-speed travel, horizontal surfaces should provide very weak (or nonexistent) laser returns. Vertical surfaces (obstacles), however, should result in stronger signals.

2 The Laser Reflectance Model

The brightness discrepancy between horizontal and vertical surfaces can be attributed to two factors, incidence angle and range. Assuming the diffuse component of the laser reflection is Lambertian, we can model the laser reflectance signal with the following relation:

$$P_{\text{return}} \propto \frac{\rho \cos \theta}{z^2}$$

where P_{return} is the power in the returning laser pulse, ρ is the actual surface reflectance ($0 \leq \rho \leq 1$), θ is the angle

of incidence of the beam with the surface and z is the range[8].

Assuming a vertical object with square corners, θ should be anywhere from zero to 45 degrees (due to rotations about the vertical z -axis). For horizontal surfaces at distances over 20m, however, θ should be between 87 and 90 degrees and should provide a much smaller return. Additionally, points on the obstacle should be closer than points on the road contained in the same scanline. In the limiting case, where the sensor beam is located at obstacle height and parallel to the road surface, laser intensity should be non-zero only on obstacle pixels.




Figure 1. A laser sensor is located at a height h on the vehicle with a depression angle α . The laser hits an obstacle of height p .

In particular, if our sensor is located at a height h , with a depression angle α (see Figure 1), assuming a flat road, we expect road pixels to return an intensity:

$$I_{\text{road}} \propto \frac{\rho_{\text{road}} \sin^3 \alpha}{h^2}$$

On the other hand, if the laser hits a vertical obstacle of height p , we would expect a returned intensity:

$$I_{\text{obs}} \propto \frac{\rho_{\text{obs}} \sin^2 \alpha \cos \alpha}{(h-p)^2}$$

Assuming the surface reflectance of the road and obstacle are comparable, i.e. $\rho_{\text{road}} \approx \rho_{\text{obs}}$ then for a typical sensor height of 1m, obstacle height of 20 cm, and sensor depression of 1° (so that the sensor is aimed at a typical look-ahead distance of approximately 60 meters away) we find that I_{obs} is nearly 2 orders of magnitude greater than I_{road} . This huge intensity discrepancy is mainly due to the difference in the incidence angle. Although this intensity discrepancy is reduced significantly for non-lambertian surfaces, it has nevertheless proven sufficient for detecting many non-lambertian obstacles.

2.1 Obstacle Tracking

Guaranteeing coverage of the road is important in an obstacle detection system. We must be sure that we are scanning frequently enough that no objects may pass undetected within our field-of-view. In fact, coverage is not enough. In a roadway environment, the detection system may occasionally detect things such as roadway ret-

roreflectors (which are highly reflective) and rain or airborne debris (which may be very near), as bright spots in the laser intensity image. To reduce the chance of labeling these bright spots as obstacles, it is important to oversample, that is to sample more frequently than necessary for guaranteeing coverage alone. The oversampling factor is the number of lines in which an obstacle of minimum height p will appear. If our vehicle is traveling at speed v , and obtaining line scans at a frequency f , then the oversampling factor n is:

$$n = \frac{pf}{v \tan \alpha}$$

For reasonable values ($v = 30$ m/s, $\alpha = 1^\circ$, $p = 0.2$ m, and $f = 36$ Hz, our current scan rate), n is approximately 14. An obstacle would appear in the line scans n times before passing beneath the sensor field-of-view at which point the object would be at a distance L , where $L = (h-p)/(\tan \alpha) = 45.8$ m.

3 Results

Recent tests with the Franklin laser scanner have shown that a laser reflectance-based system can detect 1-foot-high obstacles at ranges of 60 meters or less.

3.1 Sensor Details

All laser data in this paper was taken with the Franklin scanner, a joint effort between K²T and CMU of Pittsburgh, and Z+F of Wangen, Germany. The laser is an amplitude-modulated continuous wave (AMCW) sensor. AMCW sensors calculate distance by measuring the phase difference between the outgoing reference signal and the returned signal. The energy or amplitude of the returned signal is used as an intensity measure. The Z+F laser uses two modulation frequencies to provide 15-bit range and 16-bit intensity data. This is much greater resolution than our previous ERIM laser scanner, which provided only 8 bits of range and intensity. The improved intensity data has enabled us to detect objects further than the 20 meters possible with the ERIM[3]. For more details on the design and performance of the Z+F laser[2] and the Franklin scanner[4], readers are directed elsewhere.

3.2 Laser Intensity-Based Obstacle Detection

The image below shows a sequence of single-line laser intensity scans taken from a moving vehicle in a parking lot at CMU (see Figure 2). The intensity has been inverted for better printing. The dark areas in the inverted image (bright areas in the original image) correspond to obstacles. A schematic illustrates what happens at times t_0 , t_1 ,

and t_2 (see Figure 3).




Figure 2. A series of single-line (1-D) laser reflectance scans taken from a moving vehicle (reflectance has been scaled and inverted for better printing). Time $t=0$ is at the top of the image. Dark spots correspond to obstacles. The vehicle passes two lamp posts on the left (the dark curves), and approaches a cinder block in the center (25 to 30 meters away).




Figure 3. An overhead schematic of vehicle motion for the laser series in Figure 2 (not drawn to scale). The circular arcs represent where the laser scans would hit the ground at each of 3 time instants. The wedge at each position indicates the horizontal field of view. At time t_0 , representing the first scanline in Figure 2, the laser sweeps an arc on the ground. At time t_1 , the laser hits a lamppost, and continues to see it until time t_2 when it passes out of the field of view. At the same time, the laser hits the cinderblock near the center of its scan. Whenever the laser hits one of these objects, it gets a strong reflection (represented as a dark area in Figure 2).

Simple techniques are able to adequately separate obstacles from pavement in the image in Figure 2. A classification image based on the thresholded absolute value of a 1×5 Sobel operator is shown in Figure 4. The two long dark curves correspond to two light poles, and the dark

streak in the center corresponds to a vertical cinder block.



Figure 4. A classification of the laser reflectance scans in Figure 2, based on the absolute value of an edge operator. Black pixels indicate a potential obstacle.

In laser series in Figure 2, the intensity of the cinder block at a range of 25 to 28 meters averages 270 out of a 16-bit maximum of 65535 -- we have already lost most of our signal at this distance. Reflectance for the pavement, however, averages about 50, i.e. $I_{\text{obs}}/I_{\text{road}}$ is 5.4. Although this is an ample intensity ratio to detect the cinder block in this image, things become more difficult at longer ranges as the returned signal strength decreases. Moreover, the ratio is a far cry from the two orders of magnitude difference that was predicted. To understand why the obstacle-to-pavement intensity ratio was not as large as predicted, we needed to better characterize the reflectance model.

4 Refining the Reflectance Model

Although the first-order reflectance model and our results indicate promise for laser intensity-based obstacle detection, we need to refine the reflectance model. We are currently performing experiments to refine the model to improve our detection system and extend its range. There are two major assumptions we made in using the reflectance model to compare obstacle vs. non-obstacle laser intensities. First, we assumed that both the road surface and the obstacle were Lambertian in nature. Second, we assumed that the laser angle of incidence may be directly related to the macro geometry. In fact, the returned intensity is likely to depend on the orientation of microfacet structures in the surface. Microfacet structure and orientation may play a larger role in determining the effective angle of incidence than the macro-orientation of the surface. The large discrepancy between the predicted intensity ratio and the actual intensity ratio in the previous example suggests this may be true.

4.1 Reflectance versus Incidence Angle

To verify the reflectance model, we performed an experiment on several objects. We attached each object to a turn-

table, aimed the laser at the object, and rotated the object through at least 90 degrees while measuring laser intensity. Thousands of reflectance measurements were averaged at each angular increment to reduce the effects of sensor noise.

Intensity versus incidence angle graphs are given for two of the objects below (see Figure 5 and Figure 6). For comparison, a cosine curve of approximately equal amplitude is also shown on each figure. Intensity dropoff at shallow angles for both of these objects was similar to the behavior predicted by the Lambertian model and the macro geometry. Rougher surfaces deviate more from the predicted fall-off, though we need to perform more experiments to better characterize these deviations.

The intensity data from the first object, a piece of matte white posterboard, fits the Lambertian cosine model quite well. The intensity data for the second object, a cardboard box, does not fit the model well at incidence angles near zero. This is likely due to a specular component of the reflection.




Figure 5. Laser intensity versus incidence angle for matte white posterboard. The model fits the predicted cosine curve for a Lambertian surface.




Figure 6. Laser intensity versus incidence angle for a cardboard box. The model does not fit the predicted cosine curve for a Lambertian surface. The difference between the curves represents the energy in the specular reflection.

Since obstacle types we might encounter on the road are varied, we can not expect to build a reflectance model of any object we might encounter. Fortunately, however, specular reflections are often likely to help the system detect obstacles. Specular reflections off vertical surfaces are directed back towards the laser sensor, and depending on the size of the specular lobe may be seen by the detector. Specular reflections off horizontal surfaces, however, are directed away from the sensor. This increases our expected obstacle-to-road intensity ratio for many types of obstacles. One exception to this is obstacles with highly specular surfaces. In this case the specular energy is rarely seen by the laser sensor since it is unlikely that the laser is exactly normal to the surface. For this reason, shiny vehicle surfaces can appear dark in laser images. Even then, there is generally some portion of the vehicle which reflects strongly, allowing detection (see Figure 7). Regardless, vehicles are easily detected by a number of other sensor systems, so this is less important than being able to detect small obstacles.




Figure 7. A dark vehicle on the right of the laser image has a similar intensity to the surrounding pixels, except for very strong reflections off the headlights. The white vehicle is more easily detectable.

There were two important sources of error in this experiment. First, it was difficult to accurately measure the absolute angle of incidence. Relative angular measurements were good to within 0.1 degrees, but the initial absolute measurement was probably only accurate to within 2 degrees. Second, since the target surface could not be placed exactly at the center of rotation, the laser spot moved slightly along the surface during the rotation. Although these movements were kept within a reasonably uniform region on the surface, they would result in some intensity variations.

4.2 Reflectance versus Range

According to theory, laser intensity should be inversely proportional to the square of the range. From most of the laser images, it is clear that reflectance does fall off with range. However, some strange behavior occurs in reflectance images at low intensities. As can be seen from an image taken in the parking lot (see Figure 8), there are alternating bright and dark horizontal bands on the asphalt instead of the predicted monotonic drop in intensity as range from the vehicle increases.




Figure 8. Part of a 2-D laser reflectance image taken in a parking lot. Notice alternating bright and dark sections on the asphalt moving up the image from the bottom. (Pixel values have been scaled for better printing).

To verify the laser reflectance model and to isolate the cause of the banding in the laser intensity image, we collected single point laser data while moving a target along the optical axis. This allowed us to generate a laser intensity versus range scatterplot (see Figure 9). We found the best fit curve of the form $I = A/z^2$ where I is the laser intensity, z is the measured range, and A is a constant, and then plotted the residual between the predicted and measured intensity (see Figure 10).

At shorter distances, and consequently higher intensities, our results agree well with the model, i.e. reflectance is inversely proportional to the range squared. However, at long distances and low reflectances, the plot of the residual reveals an unmodeled oscillation in the intensity. The oscillation in the residual has a period of approximately 6.5 meters, which is the ambiguity interval of one of the laser modulation frequencies. We currently believe this is caused by optical crosstalk. A constant source of reflection in the scanner optics, such as an internal reflection from the focusing optics, would cause a reflection at a constant distance and therefore constant amplitude and phase shift. As the range to a real target is increased, the internal reflection would be alternately in phase and out of phase with the reflection from the object, which would cause the banding seen in our images[6]. We are currently conducting experiments to verify our suspicions. Although this residual is small (only 5% of the signal at 10 meters range), it becomes a significant percentage of the signal at

distances of 40 meters or more. This makes modeling and statistical-based approaches to obstacle detection more difficult.




Figure 9. Scatterplot of intensity versus range acquired by moving an object away from a non-scanning laser.




Figure 10. A scatterplot of the residual intensity error between the predicted reflectance (based on range) and the actual measured reflectance. Note the distinct oscillations in the residual. The frequency of oscillations corresponds to the ambiguity interval of the AMCW laser.

5 Conclusion

While it is unlikely that any single sensor can detect all obstacles, laser intensity-based methods appear promising for high-speed, long-distance detection of small obstacles for automated highways. Many of the obstacles we have used in our experiments so far are medium-sized (mainly for convenience in verification). However, since the detection capability is based more on angle of incidence differences between road and obstacle than size, we expect to be able to robustly detect smaller obstacles. An additional advantage of our method is that a laser intensity sensor could be lower power and far cheaper than a laser range sensor.

Nevertheless, there are likely to be some obstacles which may be very difficult to detect using laser reflectance. Very dark or highly specular surfaces may "disappear" in laser images. Laser range would not help in these cases since it is unreliable at low intensities. Hopefully, these obstacles would be detectable with radar or stereo vision.

6 Future Work

We are currently obtaining samples of asphalt and concrete. We plan to repeat the turntable and moving target experiments with these samples to better model the laser intensity for typical highway surfaces. The ease with which we can detect obstacles at long distances depends on improved modeling.

Increasing the laser power would also improve obstacle detectability at long ranges since it improves the signal-to-noise ratio. However, there are eye safety considerations so we may not be able to increase laser power much more. Since the signal-to-noise ratio depends on the laser energy reaching a point, we can also improve detection by maintaining the laser power while reducing the sampling rate, or speed, of the laser. Our current scanner has a 360 degree horizontal field-of-view. This is unnecessary and wasteful since we throw away most of the data. A better scanner design for highway-based obstacle detection would use a polygonal mirror and a smaller horizontal field-of-view. This would allow us to maintain or even increase the line rate of the sensor while decreasing the sampling rate and improving our signal-to-noise ratio. It might also be possible to reduce sensor angular resolution by a factor of two to reduce the sampling rate.

Finally, we will test our detection methods on a wider variety of obstacles and material types, including tires, wood, cement, metal, and ceramic.

7 References

- [1] Bishop, J. R. Intelligent Travel: The Automated Highway System. *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS-3)*, 1993.
- [2] Froehlich, C., M. Mettenleiter, F. Haertl. Imaging laser radar (LIDAR) for high-speed monitoring of the environment, *SPIE Proceedings on Intelligent Transportation Systems*, 1997.
- [3] Hancock, J. *High-Speed Obstacle Detection for Automated Highway Applications*. Thesis proposal. Carnegie Mellon Technical Report, CMU-RI-TR-97-17, May 1997.
- [4] Hancock, J., E. Hoffman, R. Sullivan, D. Ingimarsen, D. Langer, M. Hebert. "High-performance laser range scanner." *SPIE Proceedings on Intelligent Transportation Systems*, 1997.

- [5] Hebert, M., C. Thorpe, A. Stentz, eds. *Intelligent Unmanned Ground Vehicles*. Kluwer Academic Publishers, 1997.
- [6] Hebert, M. and E. Krotkov. 3D measurements from imaging laser radars: how good are they? *Image and Vision Computing*, vol. 10, no. 3, April 1992.
- [7] Matthies, L. and P. Grandjean. Stochastic Performance Modeling and Evaluation of Obstacle Detectability with Imaging Range Sensors. *IEEE Transactions on Robotics and Automation*, Vol. 10, No. 6, Dec. 1994.
- [8] Nitzan, D., A. Brain, and R. Duda. The Measurement and Use of Registered Reflectance and Range Data in Scene Analysis. *IEEE Proceedings*, Vol. 65, No. 2, February 1977.