

Laser Intensity-Based Obstacle Detection and Tracking

John A. Hancock

January 26, 1999

CMU-RI-TR-99-01

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy*

©1999 John A. Hancock

This research was partly sponsored by the USDOT under Cooperative Agreement Number DTFH61-94-X-00001 as part of the National Automated Highway System Consortium.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. government.

Abstract

One of the major challenges in designing intelligent vehicles capable of autonomous travel on highways is reliable obstacle detection. 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. Because of its necessity for mobile robot platforms and intelligent vehicles, there has been a great amount of research devoted to the obstacle detection problem. Although there are a number of methods that can successfully detect moving vehicles, the more difficult problem of detecting animals or small, static road debris such as tires, boxes, or crates remains unsolved.

Laser range scanners, or ladars, have been used for many years for obstacle detection. Laser scanners operate by sweeping a laser across a scene and at each angle, measuring the range and returned intensity. Past researchers have ignored the intensity signal while focusing on the range returned from the laser, since the range provides direct 3-D information useful for mapping. In this thesis, I demonstrate how laser intensity alone can be used to detect and track obstacles. While laser ranging demands fast, complicated electronics, intensity can be measured cheaply. Minimizing cost will be extremely important for any consumer system, and is a strong motivating factor for this thesis.

Laser intensity provides different information from ordinary video data since lighting and viewing directions are coincident. At long ranges and grazing angles, vertical obstacles reflect significantly more laser energy than the horizontal road. The obstacle detection system uses a high-performance laser scanner which provides fast single-line laser scans. Histogram analysis on the returned intensity signal is used to select obstacle candidates. After candidates are matched and merged with candidates from previous scans, the range to each obstacle is estimated by a novel intensity and position tracking method. Finally, the positions of all obstacles are updated based on vehicle motion before the next laser scan is acquired.

Since laser intensity has been ignored in the obstacle detection literature, I have devoted a significant portion of the thesis to examining the intensity measurements provided by the scanner. I present a laser reflectance model which provides good results for a wide variety of object surfaces. The reflectance model is based on experimental results and a combination of two popular reflectance theories from the computer graphics and computer vision literature. I also discuss road and system geometry in detail, since geometry affects the obstacle detection problem significantly.

Acknowledgments

I first want to thank my advisors, Chuck Thorpe and Martial Hébert. I learned a great deal from our work and our conversations (both academic and casual). They provided valuable insights into theory and implementation for this thesis. Thanks to the other members of my thesis committee: Eric Krotkov and Edward Riseman, whose comments and feedback were of significant help in completing the dissertation.

Thanks to Dirk Langer for his great help in maintaining several laser scanners over the last few years, and for getting me started with them in the first place. The different laser scanners I worked with were a constant source of interest (and occasional frustration). Thanks to Ryan Sullivan, who was very helpful with current scanner mechanism. Thanks also to Markus Mettenleiter at Z+F for answering many questions regarding the laser.

My officemates at CMU were a frequent source of support and interesting conversations. Thanks to Jennie Kay, Justin Boyan, Fabio Cozman, Parag Batavia, and Henry Schneiderman. Special thanks go to Parag for lending me his personal laptop for the last couple months of my thesis for data collection after a project laptop died.

Thanks to everyone on the Unmanned Ground Vehicle and Automated Highway System groups at CMU. I especially appreciate the hard work of Jim Frazier and John Kozar who kept all of our vehicles running.

My time at CMU was enriched by too many people to name them all. However, I want to thank some of my friends: Rahul and Gita Sukthankar, Mei Chen, Deepak Bapna, David LaRose, Shumeet Baluja, Todd Williamson, Rich and Cookie Madison, Yoichi and Imari Sato, Kristie Seymore, Michael Nechyba, and Stephen Chen. Thanks also to the members of the Viking Death Rats soccer team, and everyone on RoboOrg.

Finally, I want to thank my family for their support. My mother, Diane, who returned to graduate school and received her Ph.D. several years ago, inspires me with her seemingly boundless supply of energy. My father, Kenneth, who I know would be proud, always believed in me and encouraged me. My brother, Ken, and my sister, Catherine, are my confidants. My grandparents, Audrey and Kenneth Hancock and Anena Kerr have supported me both emotionally and financially.

Contents

Abstract	i
Acknowledgments	iii
Contents	v
1 Introduction	1
1.1 Thesis Overview	4
2 Obstacle Detection Methods	9
2.1 Comparison to Off-Road Obstacle Detection	9
2.2 Optical Flow	12
2.2.1 Related Work	15
2.2.2 Analysis	16
2.3 Stereo	19
2.3.1 Related Work	21
2.3.2 Analysis	23
2.4 Laser	24
2.5 Radar	26
2.6 Other Vision-Based Detection Modalities	27
2.7 Chapter Summary	27
3 System Requirements and Geometry	29
3.1 Safety Requirements	29

3.2	Road Geometry Considerations	34
3.2.1	Horizontal Road Curvature.	34
3.2.2	Vertical Road Curvature.	36
3.3	Road Geometry Impact on Laser Reflectance	39
3.4	Sensor Placement and Highway Design	41
3.5	Summary.	42
4	Scanner Modeling	45
4.1	Scanner Design	45
4.1.1	Laser Rangefinder Operational Theory	47
4.1.1.1	Pulsed Time-of-Flight.	48
4.1.1.2	Frequency-Modulated Continuous Wave (FMCW)	49
4.1.1.3	Amplitude Modulated Continuous Wave (AMCW).	51
4.1.2	Environmental Issues/Covering the Scanner	54
4.2	Performance	55
4.2.1	Mixed Pixels.	56
4.2.2	Range/Reflectance Crosstalk	58
4.2.3	Temporal Mixing	60
4.2.4	Range and Reflectance Precision	62
4.2.5	Angular Precision.	63
4.3	Summary.	65
5	Reflectance Modeling	67
5.1	The Standard Reflectance Model	68
5.2	Verification of the Standard Model	68
5.2.1	Noise Experiments	68
5.2.1.1	Photon Noise.	69
5.2.1.2	Signal Noise	72

5.2.2	Reflectance versus range	74
5.2.3	Reflectance versus incidence angle	75
5.2.4	Reflectance Models	77
5.2.4.1	Torrance and Sparrow Reflectance Model	79
5.2.4.2	Oren and Nayar Reflectance Model	81
5.2.4.3	Combining the models	82
5.2.5	Signal Noise Implications.	85
5.3	An Improved Reflectance Model	86
5.4	Intensity Drift	87
5.5	Chapter Summary/Conclusion	88
6	Obstacle Detection and Tracking	89
6.1	System Overview	89
6.1.1	Hardware	90
6.1.2	Software	90
6.2	Update.	92
6.3	Classification.	93
6.3.1	Road Model-based Thresholding	95
6.3.2	Edge-based Segmentation	96
6.3.3	Histogram-based Segmentation	97
6.4	Merge	99
6.5	Range Estimation	105
6.5.1	Least Squares	109
6.5.2	Least Median of Squares	110
6.5.3	Least Trimmed Squares	110
6.5.4	Least Squares of Maxima.	111
6.5.5	Least Squares of Medians.	112
6.5.6	Range Estimation Results.	112
6.5.7	Estimating Uncertainty in Range	115

6.6	Problems and Complications	117
6.6.1	Dark Obstacles	118
6.6.2	Ground Markings as False Positives	119
6.6.3	Optical Crosstalk	121
6.7	Chapter Summary	122
7	Experimental Results	123
7.1	Detection and Tracking Results	123
7.1.1	Indoor Experiments	124
7.1.2	Outdoor Experiments	125
7.1.2.1	Experiments without Position Information.	125
7.1.2.2	Outdoor Position-Tagged Experiments.	128
7.1.2.3	Experiments with the Environmental Cover.	130
7.1.3	Detection Results Summary	131
7.2	Range Estimation Results	132
7.2.1	Outdoor Position-Tagged Data.	132
7.2.2	Range Estimation with an Extended Field of View	135
7.2.3	Range Estimation Results Summary	141
8	Conclusion	143
8.1	Thesis Summary and Contributions	143
8.2	Future Work	145
8.2.1	Designing a Better Sensor	145
8.2.1.1	Scanner Mechanism Modifications.	145
8.2.1.2	Laser Modifications	147
8.2.2	A Radically Different Sensor	148
8.2.3	Software Improvements	149
8.2.4	Additional Testing	150

8.2.5	Active Lighting in Other Domains	151
A	A Simplified Torrance-Sparrow Model	153
B	A Kalman Filter for Obstacle Position Estimation	159
	References	163

CHAPTER 1 **Introduction**

Obstacle detection is one of the key problems in computer vision and mobile robotics. Any mobile robot that must reliably operate in an unknown or dynamic environment must be able to perform obstacle detection. It is likely that obstacle detection will never be a “solved problem.” As robots grow more capable and can operate at higher speeds, the robots will need to look farther, examine larger areas, and require smaller lag times. And as the general populace comes to depend on some of these systems, their reliability will need to be greatly enhanced. In the near future, obstacle detection for the Automated Highway System and similar projects will have to push the current limits of obstacle detection technology in many directions.

Because of its necessity for mobile robot platforms, there has been a large body of work performed on obstacle detection. Much of the obstacle detection work has been done on small laboratory robots or on vehicles navigating in natural terrain or cross-country scenarios. I will show that the high speeds involved in the highway obstacle detection problem presents difficulties that many of the methods developed for these scenarios will not overcome in the on-road scenario. Recently there has been a great amount of work for on-road

obstacle detection. As road-following systems have become more capable, more attention has been focused on the obstacle detection problem, much of it driven by programs such as the Automated Highway System (AHS) or PROMETHEUS which seek to revolutionize automobile transportation, providing consumers with a combination of “smart” cars and smart roads[4].

One way of tackling the obstacle detection problem for AHS would be to position sensors at regular intervals on the highways. The sensors would scan the road nearby for obstacles and would broadcast the location of obstacles to the local vehicles and to a maintenance center where a crew could be notified to remove the obstacle. There are two significant advantages to making the sensors part of the infrastructure. First, a sensor can use fairly simple temporal differencing techniques to find objects since it is monitoring a constant area. Second, sensor placement is more flexible; depending on the particular sensor, it could be placed very close to the ground or far above it to avoid occlusions, whichever works best for the particular sensor.

There are some significant problems associated with the deployment of a roadway infrastructure-based obstacle detection system. First, it would be extremely expensive to equip and maintain thousands of miles of highways with sensors. Second, it requires all automated vehicles to be equipped with reliable communication systems. Without “smart” vehicles to take advantage of the system, there is little incentive to build such an expensive infrastructure, and without the infrastructure, there is little incentive to buy the smart vehicles. Nevertheless, an infrastructure-based system might provide the best solution for heavily congested highways.

Because of the drawbacks of deploying an infrastructure-based system, this dissertation proposes a vehicle-based obstacle detection system that depends only on existing roadway infrastructure. The methods presented in this thesis should not be adversely affected by realistic changes to the roadway infrastructure. In addition, the detection system does not depend on a fully automated vehicle to provide benefit to the driver since it could be employed as a warning system.

Vehicle-based 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 [6],[7],[20],[25],[28],[45],[49] 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 [6],[9],[25],[28], but the more difficult problem of finding animals or 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 [4]. Although the problem of detecting static obstacles has been tackled in both the cross-country [24],[34],[41] and indoor mobile robot navigation literature [2],[5],[8] these systems have operated at low speeds (5-10 m.p.h.) and short range.

Current methods of roadway obstacle detection often fail to achieve reliable detection because the methods have not been adapted to the specific problem. Sometimes the road or sensor models (or lack thereof) are at fault -- the methods either assume too much so that the models are unrealistic, or too little so that the signal is effectively lost in the noise. Other times, the method may be poorly suited to the problem (lack of sensitivity), or basic safety requirements or complicating factors (such as unintentional sensor movement) may have been ignored. This thesis will demonstrate how small static road debris can be safely detected at long distances. In particular, it will focus on using laser reflectance. Laser reflectance, to our knowledge, has not been used in an obstacle detection system before. Laser intensity-based detection is based on a different principle of operation than most detection methods, and might be cheap enough for consumer vehicles.

Nearly all currently employed obstacle detection methods are based on direct or indirect range measurements of the surroundings. Once a range map is produced, the detection systems look for discontinuities in the range and base their detection criteria on these discontinuities. On a highway, however, a better criteria for determining whether a surface is an obstacle or not is the slope or pitch of the surface. If the surface is horizontal, it most likely belongs to the road. If the surface is vertical, however, our vehicle should avoid it.

Laser reflectance is one means of trying to obtain information about the surface orientation. The precise relationship between surface orientation and laser reflectance is discussed in Chapter 5. Polarization of reflected light may be another method to determine orientation, but is not investigated here.

Although this thesis uses an expensive scanning laser rangefinder as a testbed, an appropriate laser reflectance sensor could be built cheaply. Laser rangefinders demand complicated, fast electronics to perform range measurements. Reflection can be measured with similar optics but minimal electronics. Previous laser-based obstacle detection work has primarily used two degree-of-freedom (DOF) scanners which are expensive to build and maintain. However, adequate obstacle detection may be achieved with a fast single DOF scanning system which is likely to be cheaper and more reliable. Additionally, there is the potential for a non-scanning reflectance-based obstacle detection system based on an infrared strobe and a CCD with a narrowband optical filter.

1.1 Thesis Overview

In this dissertation, I present a methodology for laser intensity-based static obstacle detection and tracking. My goal is to demonstrate the advantages of using laser intensity as a means of on-road obstacle detection for automated vehicles over more conventional means. The experimental testbed used for much of the thesis work was Navlab 5, one of

Carnegie Mellon's automated vehicles, and a laser scanner jointly constructed by K2T of Pittsburgh, PA, and Zoller and Fröhlich of Wangen, Germany.



Figure 1-1. The testbed for the thesis work. The laser scanner built by K2T and Zoller and Fröhlich is mounted on the roof of Navlab 5, one of Carnegie Mellon's automated vehicles.

In Chapter 2, I examine previous obstacle detection work. An analytical comparison of optical flow shows that it is unlikely to detect small static obstacles at long ranges. In contrast, I show that stereo vision should be capable of solving the detection problem. Stereo vision, however, requires significant computing resources and does not operate well at night. Laser range methods are discussed and show promise. Laser-based systems work at night (better than during the daytime), but scanning laser rangefinders currently are too expensive, and adequate eye-safe power is a concern. Radar is an excellent means of detecting vehicles, but fails to find small obstacles. There is no current obstacle detection method that works in all environments with the necessary reliability.

When work on intelligent vehicles first began, most efforts were focused on road-following. As road-following solutions improved, obstacle detection systems were built to detect other vehicles. As detection systems improved, researchers began to refocus their efforts on detecting smaller objects. Throughout these changes, sensor placement often remained the same (on top of the vehicle). During the switch from vehicle detection to

small objects, systems often used the same methods and assumptions. Few researchers examined the obstacle detection system as a whole or verified the assumptions concerning road geometry inherent in their methods. Chapter 3 tries to fill these gaps. I examine the constraints placed on the obstacle detection system by vehicle and road parameters such as vehicle velocity and road curvature. I provide analyses to determine required and/or optimal parameters for system throughput, field of view, sensor location, and other system parameters. Actual highway design parameters are used to show that the flat-world model used in many highway obstacle detection algorithms is invalid for detecting small objects on the basis of height.

Chapter 4 discusses the operational theory, design, and performance of our current laser scanning system. The scanner uses a single mirror with 2 degrees-of-freedom (DOF) to deflect a laser beam around a scene with a 360 degree horizontal field-of-view and a 30 degree vertical field-of-view. At each instant, the scanning mechanism records the position of the mirror and the range and intensity measured by the laser. I describe the 3 major laser ranging methodologies: pulsed time-of-flight (TOF), amplitude-modulated continuous wave (AMCW), and frequency-modulated continuous wave (FMCW). Problems particular to our method of ranging (phase differencing an AMCW laser signal), are discussed in detail. One problem in the design that was somewhat unexpected was the difficulty in designing a suitable optical protective cover for the laser; internal reflections from the cover cause significant problems. Finally, Chapter 4 also provides results of scanner precision tests.

Since laser intensity has been essentially ignored by researchers, it was important to examine and verify current laser reflection models. The only reflection model provided by the laser scanner literature assumes Lambertian surfaces and fails to characterize noise in the intensity. As any computer vision researcher knows, most surfaces are not Lambertian. Chapter 5 describes experiments performed to analyze intensity noise characteristics and the relationship of laser intensity to surface range and incidence angle. Two reflection models from the computer graphics and computer vision literature are used to help build an improved laser reflectance model.

Chapter 6 details the obstacle detection and tracking system. The system hardware consists of the laser scanner (as described in Chapter 4) and a vehicle position sensor. Vertical nodding on the scanner is disabled so that the scanner collects single laser lines at constant elevation. The software algorithm is divided into 4 steps. A segmentation module classifies pixels in a new laser scanline as obstacle or non-obstacle based on the laser intensity. Obstacle candidates from this scanline are then matched and merged with obstacles from previous scans if possible. A memory-based algorithm is used to estimate the range to each merged obstacle using the history of laser intensity measurements and vehicle position. Before the next scan is taken, the vehicle-relative positions of all obstacles are updated based on a vehicle motion estimate. The system is capable of detecting and tracking a wide variety of obstacles based solely on the intensity.

Chapter 7 provides further experimental results in obstacle detection, tracking, and range estimation. In most experiments, the system succeeds in obstacle detection without any false positives. Intensity-based range estimation is shown to converge to the correct range for most diffuse surfaces.

Finally, Chapter 8 reviews the system capabilities and points of failure. The thesis demonstrates the potential of using laser intensity as an obstacle detection method. I discuss the thesis contributions in system analysis, physical modeling, and computer vision methods. I provide design suggestions for an improved laser intensity sensor, and conclude with several ideas for future work.

CHAPTER 2

Obstacle Detection Methods

There are a variety of competing methods that have been proposed for obstacle detection and many papers written on such methods. There have been few papers, however, which have offered more than a trivial comparison of the various methods, especially in any analytical fashion. This chapter seeks to fill that gap by examining the various obstacle detection methods for a highway environment from an analytic point of view. By examining each of the basic methods, I hope to illuminate which methods ought to work best for the specific problem of on-road obstacle detection at highway speeds.

2.1 Comparison to Off-Road Obstacle Detection

There has also been extensive work on the off-road, or cross-country, obstacle detection problem. Two natural questions to ask are: Can we learn anything from this body of work? Can we use the same methods? The answers are “yes” and “no.” We may categorize the requirements for an obstacle detection system into two basic types: system requirements and detectability/sensor requirements. The system requirements try to abstract from the

specific sensor and deal with response and throughput of the system. System requirements and system geometry are discussed further in Chapter 3. The detectability requirements deal with the accuracy and resolution of the sensor and are often at odds with the system requirements. A quick comparison of the off-road and on-road problems will show that we cannot use the off-road detection methods without serious modification and still hope to succeed. However, much of the basic system analysis still holds.

Basic system analysis has been neglected from most obstacle detection papers, but it is necessary to examine both the detection requirements and the system requirements as a whole so that adequate trade-offs can be made to ensure good system performance. This chapter will explore several basic methods used for on-road obstacle detection including optical flow and stereo, but I will also briefly consider the possibility of other detection modalities.

Through no lack of effort, there have been no satisfactory solutions so far to the problem of small static obstacle detection at highway speeds. Although there have been good results published for vehicle detection, there have been few results reported on the detection of small static obstacles. Moreover, these few results have generally been reported in vague terms which give the reader little ability to compare methods. Papers on cross-country obstacle detection systems are typically no better at reporting results in a standardized way, although this is probably more excusable since rough terrain is difficult to describe in an analytical fashion and an obstacle may be less well-defined in cross-country applications. There has been some success in cross-country obstacle detection, so it is worth briefly examining the problem.

Typical speeds for cross-country applications are on the order of 5 to 10 m.p.h. Typical highway speeds, however, are many times this. Since stopping distance is proportional to the square of the speed, the stopping distance for on-road applications is much larger, perhaps by 2 orders of magnitude. In addition, a cross-country vehicle moving at relatively slow speeds can climb larger objects without damage to the vehicle than the typical passenger car travelling on the highway. For the highway scenario, we must be able to detect any

objects larger than a few inches. Coupling these facts together, we see that the resolution necessary for on-road obstacle detection may be 2 to 3 orders of magnitude greater if we use standard techniques from cross-country with similar field-of-view. Given higher speeds, we must also examine a greater amount of terrain. Although the effect of latencies on lookahead distance are generally dwarfed by the stopping distance for highway applications, it is still important to have small latencies in the processing system. High latency systems can cause problems in vehicle control.

These difficulties might make the on-road problem seem almost intractable given that the off-road obstacle detection problem is still far from being solved. Fortunately, however, roads are reasonably planar. Given our previous analysis, it should be quite clear that we need to use this fact in order to win in this scenario. Assuming planarity of the road improves our signal-to-noise ratio greatly. This simplifies a number of things. We don't need to build maps, estimate heights of objects, and check for rollover conditions, etc. We only need to find areas which violate the planarity constraint and avoid them.

For many of the computer vision methods I describe, there are both dense and discrete methods. Dense methods provide some estimate (range, optical flow, etc.) at every image pixel. Discrete methods first use some feature detection method to pick out candidate potential obstacles and then attempt to calculate range, optical flow, or some other measure in order to determine whether a feature (or group of features) actually belongs to an obstacle. In the near future, it is likely that only discrete methods will be able to achieve the low processing latencies necessary for highway obstacle detection without special purpose hardware. For this reason, this chapter focuses on discrete methods, although I do not ignore the dense methods since these will become feasible as computers become faster.

I name the process of finding obstacles once given a set of image features or patches, obstacle determination. I focus on the problem of obstacle determination in this chapter. Because an image is a mapping of 3-D space to 2-D space, a single image can not determine whether an image patch belongs to the ground plane or not (see Figure 2-1). Additional information, in the form of more images or certain assumptions or a priori knowledge about

potential obstacles, must be used to perform obstacle determination. If an image feature or patch belongs to the ground plane, then it is not a threat. If it does not, however, it should be considered an obstacle and avoided. I refer to the obstacle determination problem throughout the chapter.




Figure 2-1. A single image cannot distinguish whether image patch C is caused by A (an obstacle), or B (a ground plane feature, i.e. no obstacle). We call the problem of using additional information to disambiguate the two obstacle determination.

2.2 Optical Flow

One method of performing obstacle determination is optical flow. Optical flow uses the motion of either the objects or the camera to determine the structure of the environment. Optical flow uses two or more images taken at different times to perform obstacle detection. The basic concept of using optical flow for obstacle determination is illustrated below (see Figure 2-2 on page 13). A point in the image, y_0 , can be mapped to either obstacle point P or ground point Q (or any point in between) at time t_0 . At time t_1 , the camera takes another image after the vehicle has moved forward by distance d . A matching technique or flow calculation is employed to map y_0 to a point in the new image. If this point is close to y_{obs} , then it is likely that it corresponds to an obstacle. If it is closer to y_{gnd} , however, then it most likely corresponds to a ground-plane point. By examining the flow (in the y -direction in the illustrated case), i.e. the movement of y_0 at t_0 to a new point in the second image,

the algorithm can theoretically detect whether the point belongs to an obstacle or to the ground plane.

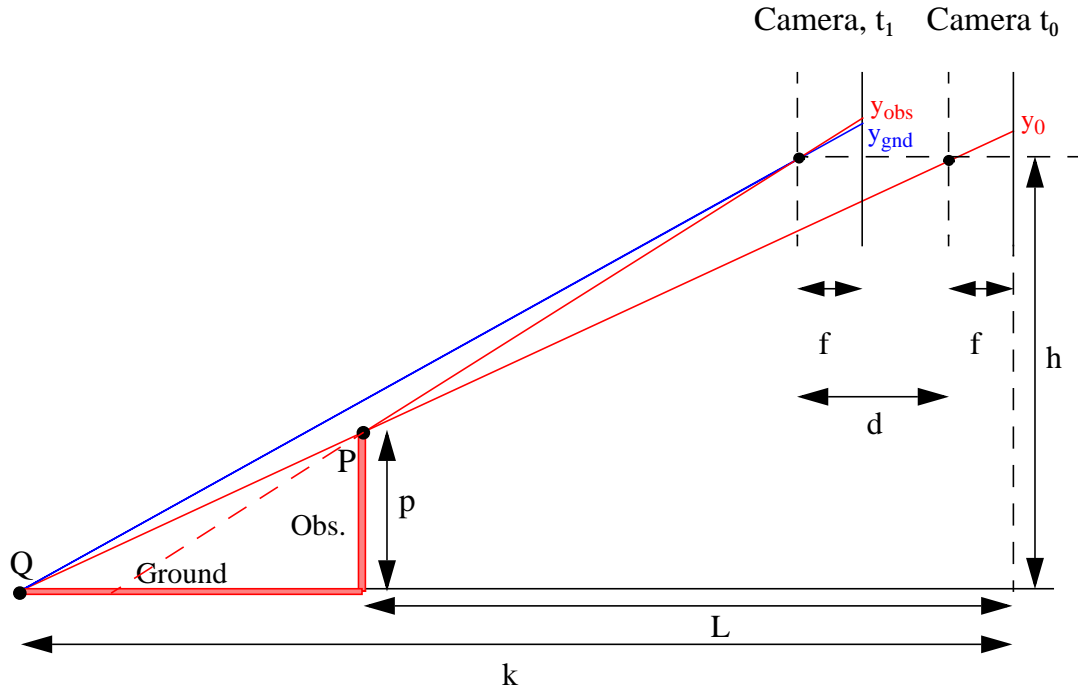


Figure 2-2. A point y_0 on the image plane at time t_0 can be mapped to either point P or Q in the world (or any point in-between). Another image taken at time t_1 after the camera has moved a distance d can be used to disambiguate between P and Q depending on whether it is mapped to y_{obs} or y_{gnd} .

Enkelmann describes this as computing a model flow vector field for the image assuming a flat world, and comparing the magnitudes of the model flow vector and the measured flow vector at the top of the obstacle[11]. There have been a number of optical flow or similar techniques proposed in the obstacle detection literature. Although some of these techniques have had success in detecting large static obstacles, optical flow methods are best at detecting moving objects which have motion fields significantly different from the background (caused by the egomotion of the vehicle). Another vehicle, for example, can be detected and tracked using optical flow techniques, even if it is moving with low relative velocity with respect to our own vehicle because it is still moving with respect to the background (the road). This may be sensed by comparing the direction of the calculated flow vector with the direction of the model flow vector. I briefly describe several methods for calculating optical flow, and then illustrate why optical flow is too insensitive to be effec-

tively used in the basic static obstacle determination problem for small obstacles at large distances.

The three basic methods for computing optical flow for obstacle detection were described by Enkelmann: analytical, edge-based, and discrete.

Analytical methods generally first perform spatio-temporal smoothing over a sequence of images. The algorithms then use the optical flow equation:

$$E_x u + E_y v + E_t = 0$$

Equation 2-1. The optical flow equation. E_x , E_y , and E_t are the estimated spatial and temporal derivatives of intensity at a pixel, and u and v are the x and y components of the optical flow.

Since a single equation is not enough to solve for two unknown flow components, smoothness is assumed to provide a second constraint. Although analytical methods often generate the best and densest flow maps, they are very computationally expensive and may not work for small objects as the spatio-temporal filtering and the smoothness constraint will tend to blend them into the background. The spatio-temporal filtering has the additional effect (if done symmetrically about the current time) of magnifying the effective sensor acquisition latency by forcing the processing to wait for future images to be captured before it can calculate the flow for the current image.

Edge-based methods have been used with success for detecting vehicles and tracking them over time[6],[11]. Typically, these methods involve an edge extraction step followed by a dynamic programming method to match edges between images. The algorithms tend to run faster than analytical methods since they compute one flow vector only for every edge rather than every pixel. Reliable edges can be matched fairly robustly in these schemes because of the length and orientation constraints that can be used in matching. One problem with these methods is that the flow component along the edge is generally unreliable insofar as the edge-detectors themselves often find broken lines, etc. Broken lines can cause mismatches that make the flow component normal to the edge wrong too, but this happens less frequently.

Pixel-based discrete methods, on the other hand, only compute the optical flow at points where the flow is fully constrained without resorting to smoothness constraints. Typically, the algorithm will choose local intensity minima or maxima or use an interest operator, e.g. Moravec interest operator, to select points in an image. Then the algorithm matches points using correlation similar to that in stereo matching. Using the observation that under a planar ground assumption, an obstacle only changes the magnitude and not the direction of the flow vector, it is possible to constrain the search to a single dimension along the direction of the model flow vector (with the drawback that it may not allow the detection of moving objects). The flow vector for the pixel is then just the vector from the original pixel position to the location of the maximum correlation in the new image.

2.2.1 Related Work

An example of recent work that uses the analytic approach to optical flow is that of Kruger, et. al. Their system uses custom image processing hardware to compute the necessary spatiotemporal derivatives[28]. The optical flow vectors are then clustered to eliminate outliers and to reduce the amount of data. The aggregate flow vectors are then compared against the model flow vectors. If the calculated flow vector doesn't match the model vector, the cluster is considered a moving object. A second test is performed to isolate stationary obstacle candidates.

A related technique is one proposed by Heisele and Ritter which uses color blob flow to detect obstacles[20]. Rather than clustering in flow vector space, however, Heisele and Ritter use clustering in color space to segment the image and obtain a symbolic description of the color blobs in the image. The blobs are then tracked over images using several heuristics including consistency in color, area, and aspect ratio of the bounding box. A single flow vector can then be calculated for the blob as a whole. Different blobs can be grouped into a single motion segment if they have similar centroid motions.

In recent years, there have also been a number of papers proposing methods that only calculate one component of the optical flow, rather than both.

Ancona described a static obstacle detection method based on optical flow[2]. The method was developed for an indoor mobile robot, so the primary assumption of a flat ground plane still applies. It is similar to the feature-based optical flow method in that they estimate the optical flow using a correlation scheme. However, it is computed on only one row of the image plane and the methodology is different in that it is the temporal evolution of the predicted velocity field that is used to detect an obstacle (rather than a direct comparison of the expected model flow vector versus the calculated flow). For this reason, no calibration is necessary. Ancona uses the same cue as Enkelmann for detecting static obstacles: in translatory motion, an obstacle does not change the direction of the vector flow, but only increases its length. However, it is shown from the example, that the method is fairly insensitive, and the robot is only able to detect an obstacle once it is fairly close, which is clearly inadequate for a high-speed vehicle.

Bohrer, et. al., also describe an optical flow method for an indoor mobile robot[5]. Bohrer warps his images using a “inverse perspective mapping” (IPM) to project the image onto the ground plane. Obstacles are detected by comparing the magnitude of the actual flow vector to the model vector obtained from the IPM. Again, however, the method is used at relatively short ranges (0.5 to 2.5 m) and large camera tilt angles which greatly improves the sensitivity of optical flow.

2.2.2 Analysis

I now show that optical flow is too insensitive for small static obstacle detection at long distances. I use a slightly different metric for the optical flow which operates on image patches rather than individual pixels (see Figure 2-3 on page 17) for several reasons. First, I expect to use discrete methods that will first detect anomalous image patches so it seems natural to compute the flow on the patch itself. Second, using image patches simplifies some of the relationships in the equations. Third, it requires less calibration. Finally, it will

help us to illustrate a point about the relationship of the area of an image patch to its optical flow.

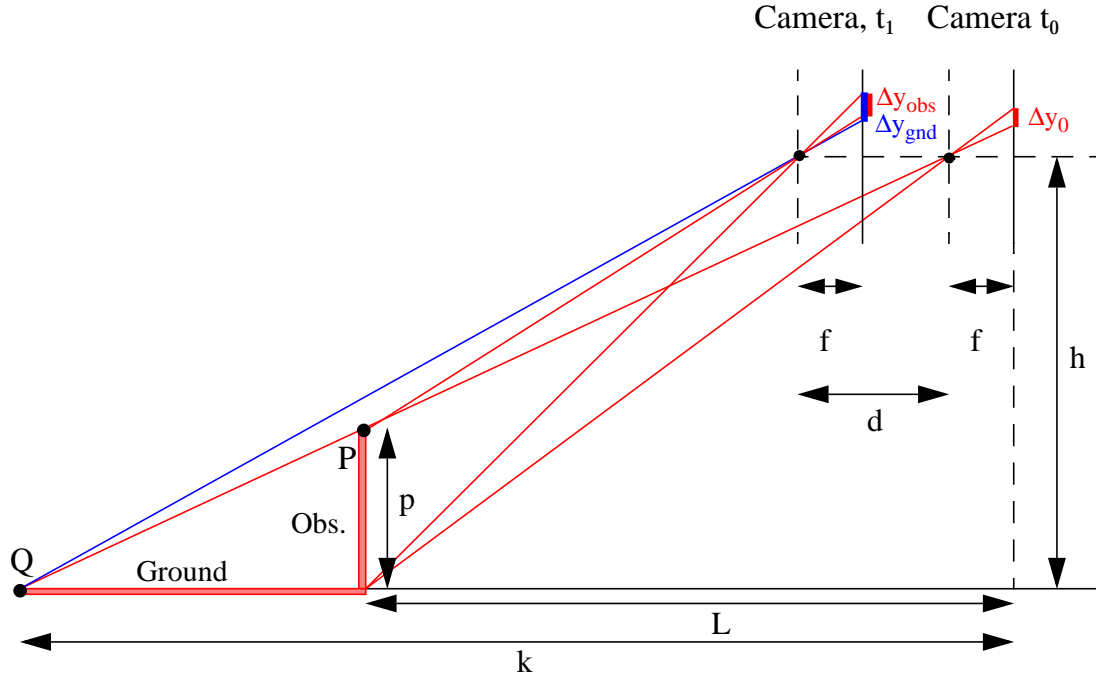


Figure 2-3. Optical flow can reveal whether the object seen in the camera at time t_0 (in red), is actually an obstacle (Obs.) of height p , or just a ground plane discoloration (Ground). The edges of the object are separated by Δy_0 . At time t_1 , after the camera has moved a distance d towards the object, the camera will see the area marked in red if the image patch at t_0 was an obstacle with the edges separated by Δy_{obs} . If not an obstacle, the area will appear larger in the camera, consisting of the blue patch, with the edges separated by Δy_{gnd} .

From the diagram above, at t_0 , the horizontal edges of the image patch are separated by Δy_0 . At time t_1 , the horizontal edges of the image patch are separated by Δy_{obs} if the patch corresponds to an obstacle, and Δy_{gnd} if not. Then we have the following equations:

$$\frac{\Delta y_0}{f} = \frac{p}{L} \quad \frac{\Delta y_{obs}}{f} = \frac{p}{L-d} \quad \frac{\Delta y_{gnd}}{f} = h \left(\frac{1}{L-d} - \frac{1}{k-d} \right)$$

$$k = L \frac{h}{h-p}$$

To keep the equations simple, I have assumed a horizontally-oriented camera and not included the tilt of the camera in the equations. For typical values for a highway scenario,

we find that the tilt of the camera makes very little difference in the values we are calculating (around 0.1%). The percentage difference between Δy_{obs} and Δy_{gnd} is then:

$$\frac{\Delta y_{gnd} - \Delta y_{obs}}{\Delta y_{gnd}} = \frac{d(h - p)}{hL}$$

Equation 2-2. The percentage difference between the height of the patch in the image is determined by the temporal baseline d , the camera height h , the obstacle height p , and the target distance L .

This is nearly equivalent to the measure that we obtain if we compare the magnitudes of the model flow vector, and the measured flow vector. However, this result allows us to see the dependencies on the variables a little more easily, so we will use it. For h greater than p , it is clear that our obstacle determination measure will not be larger than d/L , i.e. the temporal baseline divided by the object distance.

Chapter 3 provides a discussion of the geometry and provides typical values for the above variables for a highway scenario. If we substitute the typical values ($d = 5$ meters, $L = 60$ meters, $h = 1.0$ meters, $p = 0.2$ m) into the above equations, we find that there is approximately a 7% difference in the growth of Δy_{obs} and Δy_{gnd} . This is too small a difference to robustly detect in general because optical flow calculations are sensitive to noise. Enkelmann reported that he used a threshold of 30% in the change in magnitude of the flow vector (which is a nearly equivalent measure) to detect an obstacle[11].

By increasing the temporal baseline, d , we can improve the object's detectability but at a cost -- we need to push our lookahead distance back by just as much. Increasing our temporal baseline also has the drawbacks of increasing the effective sensor latency and possibly increasing memory requirements as we may need to maintain a queue of back images for processing until our temporal baseline is long enough for the current image.

Since real images are discrete, we should check our flow metric using pixel measurements and some typical values for a CCD camera. To obtain an n pixel difference between Δy_{obs} and Δy_{gnd} would require a focal length of:

$$f = nc \frac{(d - L)(dh - hL - dp)}{dp(h - p)}$$

Equation 2-3. To detect an n pixel difference in the vertical flow would require a focal length f based on the temporal baseline d , camera height h , obstacle height p , obstacle distance L , and the CCD pixel size c .

This amounts to requiring a 69 mm lens to obtain a 2 pixel difference on a small (1/2") CCD camera and corresponds to approximately a 6 degree field of view. This is a very long focal length for a CCD camera. A sub-pixel correlation method or an analytic approach might achieve accurate sub-pixel measurements which could reduce the necessary lens focal length.

Since images are discrete, it is sometimes useful to use image patch areas as a metric rather than segment lengths since calculations using areas, which rely on more pixels, are less sensitive to a noisy pixel or two. To show how optical flow relates to the area of the image patch, we can calculate Δx_{obs} and Δx_{gnd} (along the top of the obstacle of width w) to get the ratio:

$$\frac{\Delta x_{gnd} - \Delta x_{obs}}{\Delta x_{gnd}} = \frac{dp}{h(d - L)}$$

Equation 2-4. The percentage difference in the width of the image patch is determined by the temporal baseline d , camera height h , obstacle height p , and obstacle distance L .

Note that since $d < L$, this change is of opposite sign (and several times smaller for typical numbers) from the ratio using the y-coordinates. This indicates that, in general, comparing the measured area of the image patch, $\Delta y_{obs}\Delta x_{obs}$, to the predicted area $\Delta y_{gnd}\Delta x_{gnd}$, would be a worse indicator for obstacle determination than just image patch height alone. On the other hand, using the aspect ratio would improve sensitivity slightly (there is approximately a 10% difference between the obstacle and non-obstacle aspect ratios at time t_1). Regardless of the slight improvement obtainable by using aspect ratio, none of the optical flow methods seem likely to be able to detect small (20 cm) static obstacles at 60 meters without subpixel accuracy and a very long focal length lens.

2.3 Stereo

Binocular stereo vision has been a popular approach for obstacle detection in both cross-country and on-road scenarios. Stereo can be calculated in either a dense or discrete (point or edge-based) manner, but its calculation is the same in all cases. Discrete stereo vision is very similar to discrete optical flow except that a spatial baseline is used rather

than a temporal one. For point-based discrete stereo methods, an identical approach can be used to choose the features for correspondence as in optical flow, i.e. an operator which finds local minima, maxima, or corners. A correlation approach is then used to perform correspondence between the features in the first image and features in the second image (hopefully the same). In edge-based methods, the same edges that are matched in optical flow calculations can be used for matching in stereo algorithms.

As in optical flow, the search space for a corresponding point in the second image can be reduced to one dimension. In the case of stereo, the match must fall along the epipolar line. Given a single image, the same features are found and essentially the same matching procedure is run to match features to the second image, although the search is made along different directions. Stereo has the advantage, however, that if the cameras are situated with the optical axes perpendicular to the line connecting the optical centers, the search space is constrained to the same image row. This makes stereo simpler to implement and faster than optical flow. The conceptual difference is that one set of images is taken with a spatial base-

line, and one with a temporal baseline. I show that in these highway scenarios, the spatial baseline provides better detectability.




Figure 2-4. Stereo camera setup. The distance of the point from the cameras is proportional to the baseline and inversely proportional to the disparity.

The basic stereo equation is as follows:

$$z = \frac{bf}{x_2 - x_1} = \frac{bf}{d}$$

Equation 2-5. The stereo equation for the depth z of a point. b is the length of the camera baseline, f is the camera focal length, and d is the disparity between the x -location of the feature in each image.

2.3.1 Related Work

A number of systems described in the intelligent vehicle literature use stereo as an obstacle detection method.

Successful stereo algorithms can be quite complicated. JPL has built a real-time stereo vision system for autonomous cross-country navigation[35]. An eleven-step process is described by Matthies, et. al. The major steps include computation of image pyramids, and a computation of the SSD for 7×7 windows over a fixed disparity range. The minimum SSD correlation for each pixel is used to calculate the disparity for each pixel. Bad matches

are then filtered by using a left-right-line-of-sight (LRLOS) consistency check. Sub-pixel disparity is estimated by fitting parabolas to the SSD values to estimate the location of the true SSD minimum. The disparity map is then smoothed to reduce noise. Small regions are filtered out to eliminate bad matches. Triangulation is then used to produce X-Y-Z coordinates for each pixel and then transformed to the vehicle coordinate frame. Positive obstacles are then finally detected by thresholding the output of a slope operator to the range image. This large number of steps is necessary for a reliable dense depth map suitable for cross-country navigation. However, simpler methods can be used for on-road stereo detection.

Reliable sub-pixel disparity calculation has been investigated by Kaminski, et. al.[22].The authors point out that reliable sub-pixel offers the potential advantage that stereo systems could be built with smaller baselines, allowing them to potentially be packaged into a rear view mirror.

As in his optical flow work, Bohrer uses an inverse-perspective method in his stereo vision system for highway obstacle detection[6]. In this paradigm, the disparities are not directly dealt with to detect an obstacle. Instead, the left and right images are warped through an inverse-perspective mapping so that all ground plane points have zero disparity. Simple differencing or low correlation values between identically-located image points corresponds to an obstacle. This sort of method works especially well in images with little texture (such as roads) because the algorithm is not forced to choose a best match for a pixel in an area where all correlation values may be very close, but only must determine whether the correlation value is high enough. Bohrer, et. al. describe a stereo system, Vision-Bumper, which uses an inverse-perspective paradigm for detecting other vehicles in the 5 to 35 meter range. The system uses 1/2" CCD cameras and a 1.16m baseline.

There has also been some work done in performing highway obstacle detection with linear cameras and stereo [7], [45]. Linear cameras have the advantage that they reduce the overall amount of data to be processed, but increase the linear resolution which gives them a better sensitivity for depth calculations. One drawback, however, is that matching can be

more difficult and less reliable with only one dimension of information. Alternative methods have been designed for both edge detection and matching within the linear stereo images. Linear cameras are also more difficult to calibrate, but Bruyelle and Postaire have developed a method for this as well. Ruichek and Postaire proposed a neural approach to performing linear stereo matching. Their method uses a Hopfield net to perform an optimization based on an energy function instead of standard correlation. Although the authors report that they are able to detect a person at 50 meters, this unfortunately means very little analytically, since we are not told at what height the sensor is aimed (a person is a very tall obstacle).

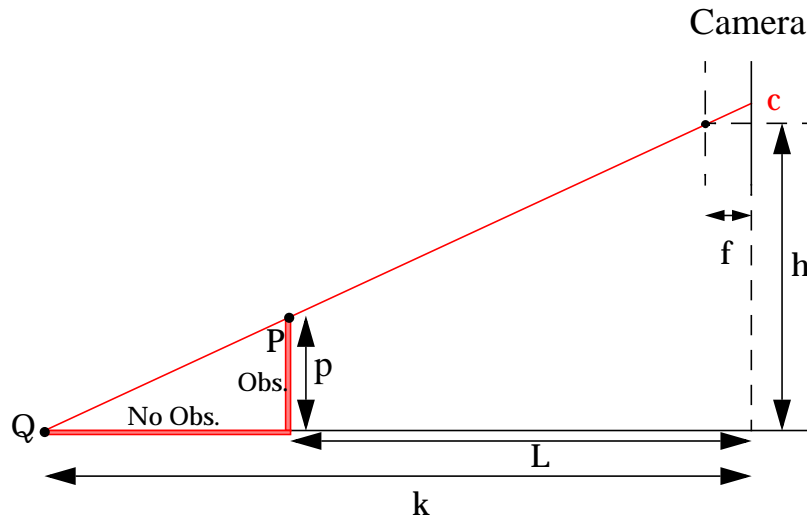


Figure 2-5. A point c in an image may correspond to either obstacle point P or ground point Q in the world.

2.3.2 Analysis

Given a setup similar to that illustrated for the optical flow, with an obstacle of height p located at a distance L , image point c in each of the two stereo cameras will image the point P if there is an obstacle and the point Q if there is no obstacle. The percentage difference in the calculated disparity for these two cases is:

$$\frac{z_{noobs} - z_{obs}}{z_{noobs}} = \frac{k - L}{k} = \frac{p}{h}$$

Equation 2-6. z_{obs} is the range to the top of an obstacle of height p . z_{noobs} is the range to the road behind it from the camera point-of-view at height h .

If we consider the discrete nature of the image to calculate the focal length necessary to result in an n -pixel change in disparity (as we did with optical flow earlier), we find:

$$f = \frac{ncLh}{bp}$$

Equation 2-7. A focal length f is needed to detect an n pixel disparity given stereo baseline b , camera height h , obstacle height p , obstacle distance L , and CCD pixel size c .

Using the same geometric parameters as in the optical flow case, we compare the sensitivity of stereo to optical flow. For a 2-pixel disparity and a baseline, $b = 1$ m, we find that $f = 9.5$ mm, a quite reasonable focal length (much shorter than that required for optical flow).

With a well-constructed stereo jig, stereo has the further advantage that both cameras will have a similar roll and pitch. In optical flow, by using a temporal baseline, we have no such guarantee -- the pitch and roll of the vehicle and camera system are likely to change somewhat from image to image. This makes matching and accurate flow vector estimation even more difficult. Bohrer mentions that a car typically has pitch angle variations of +/- 5 degrees [6].

2.4 Laser

Although we are unaware of any systems that use laser reflectance for obstacle detection, there have been many systems which have used laser range measurements for obstacle detection. Range measurements can be either indirect or direct.

Laser striping is an indirect ranging method that has been used on a number of robots (especially indoor robots). Laser striping uses a fixed CCD camera and a laser stripe (visible to the CCD) that is swept across a scene. Triangulating between the known direction of the laser and the direction provided by the laser stripe position on the CCD provides a range estimate at each image row. Laser striping can provide 3-D geometry fairly quickly since the laser only needs to be scanned in one direction, and computation is simple. There are a couple problems with laser striping. First, the laser must be easily detectable within the image (significantly stronger than ambient light). As with any triangulation system, range accuracy improves as the distance between the laser and CCD is increased. However, as

this distance is increased, the problem of “shadowing” worsens. Shadowing occurs when points visible to the laser may not be visible to the camera and vice versa. For the simplest systems (a beam diverged by a cylindrical lens), laser striping requires more laser power than direct methods and hence is most useful indoors or at short distances.

Laser rangefinders avoid the shadowing problem by keeping the transmitted and received beams approximately coaxial, and measuring range directly. A 3-D laser scanner operates by sweeping a laser across the scene in two dimensions. At each pixel, the instrument measures the time that it takes for a laser beam to leave the sensor, strike a surface, and return. There are several methods for measuring the time, and these are described in Chapter 4. Many sensors also provide an intensity measurement at each pixel by measuring the energy of the returned laser signal. Thus, a full sweep in two-dimensions can provide both a depth map and an intensity image.

Once the laser provides a depth map for the scene, the data is typically transformed into an elevation map based on the known position of the sensor. Terrain-typing algorithms may then be run on the resulting elevation map, and discrete obstacles may be detected by looking for discontinuities in the elevation map. However, there are a number of problems with this method. First, it is computationally expensive. Second, the discrete grid size cannot represent vertical surfaces, and hence may miss some potentially dangerous obstacles.

Most ranging software systems have been image-based, meaning that the system waited for an entire laser image to be acquired and processed before obstacles were reported to the path generation module. Hebert proposed a system that processes range data pixel by pixel to reduce the latencies involved and improve system efficiency[18]. A pixel by pixel method also reduces the dependency on a particular range sensor, since methods which use entire images are tuned to the specific field-of-view and geometry of the sensor. Reported latencies were reduced to under 100 ms or under. With the addition of a planning module, the system was demonstrated in multi-kilometer traverses through unknown terrain[32].

Despite problems with previous laser range methods, laser range would likely provide adequate obstacle detection capability for highway environments given adequate laser

power. The distance between an obstacle and the road surface behind it is the same as in the stereo vision system described in Equation 2-6. A laser sensor mounted at a height of 1.5 meters looking at the road at 60 meters would see an obstacle of 20 cm at a range of 52 meters. Given adequate power, this difference should be easily detectable by a laser rangefinder. Our current laser scanner is not powerful enough to provide reliable range estimates from the road at these distances. The real problem with laser rangefinders currently is that they are too expensive as a solution for obstacle detection for consumer vehicles.

Although laser reflectance has been ignored by many researchers, there has been some work on characterizing laser reflectance. In particular, Nitzan, Brain, and Duda, provided an analysis of laser range and reflectance data with an early laser rangefinder, including signal average and noise characteristics[37]. They also showed intensity information could be normalized for range effects. Eberle performed some early work on reflectance processing at ERIM with their laser scanner (and the first laser scanner used at CMU). In particular, Eberle also normalized the reflectance for the effects of range, and showed that it could detect some obstacles[10].

2.5 Radar

Radar is an excellent means of detecting other vehicles because radar works at long ranges and is relatively unaffected by rain or snow. One such radar for vehicle detection was detailed in Langer's Ph.D. thesis[30]. The radar was capable of detecting vehicles at distances of up to 200 meters with a range resolution of approximately 0.1 meters. The sensor had a 3° vertical field of view and a 12° horizontal field of view. Bearing to a target could be estimated via wavefront reconstruction, and when combined with geometric information about the road, potential obstacles could be mapped to an individual lane. Since radars provide direct range and may also provide a doppler velocity measurement, they will most likely be a standard sensor for automated vehicles.

Unfortunately, current radars are not able to reliably detect small objects at ample distances. Metal surfaces are good radar reflectors, and hence make vehicle detection fairly easy. The ease with which an object may be detected at a given range is related to its *radar*

cross section. Vehicles have a much larger radar cross section (10 m^2) than people (0.2 to 2 m^2)[30], and most road debris will have an even smaller radar cross section, making them undetectable.

2.6 Other Vision-Based Detection Modalities

Although the most promising obstacle detection methods have been discussed, there is the potential for new methods based on numerous visual cues. At least 12 depth cues are available to humans, only two of which are binocular. And there are other cues unrelated to depth that are used for obstacle detection (shape, color, luminance, etc.)[26]. Unfortunately, many of the depth cues operate at distances shorter or longer than the distances with which we are concerned (on the order of 50 meters). Most are also difficult to encode in a computer program because they require a priori knowledge: typical size of the obstacle, etc. It would be impractical to teach a system what every potential obstacle might look like. Nevertheless, some of these other cues may be useful to future obstacle detection systems.

One depth cue that can easily be used on highways to estimate the distance to a suspected obstacle, however, is linear perspective effects. Linear perspective states that the distance between two points subtends a smaller angle at larger depths. Since we know approximately the width of the road, we can use the linear perspective effects to estimate the distance to a given image row. Relative height may be used as a depth cue because given level ground, objects which are closer to the horizon are also farther away. Once calibrated using linear perspective effects, the vertical location of an object in the image could be used to estimate its distance.

2.7 Chapter Summary

Currently there are no satisfying solutions to the obstacle detection problem. All systems have their drawbacks, and almost all methods are susceptible to problems caused by environmental conditions such as ambient lighting (or lack thereof) or precipitation.

The analyses in Section 2.2.2 and Section 2.3.2 show that a vehicle-mounted optical flow system is unlikely to be successful for static obstacle detection, but that a stereo system is likely feasible for obstacle detection in daylight.

Current laser range-based methods have not been proven at highway distances and are too expensive as a solution for consumer vehicle obstacle detection. Laser intensity has been ignored as a potential solution by most researchers, and this thesis fills that gap by examining laser scanners and laser intensity in more detail in later chapters.

Radar, while an excellent means of detecting vehicles, is an unlikely candidate for detecting small static obstacles because of their relatively low radar cross section.

It is likely that multiple detection methods will be necessary for reliability. For example, the combination of stereo and laser shows promise since stereo works well in daylight, and a laser system works best in little or no ambient light.

The next chapter provides a more detailed discussion of the system geometry that was used in the analyses in this chapter.

CHAPTER 3

System Requirements and Geometry

Basic system analysis has been neglected from most obstacle detection papers, but it is necessary to examine both the detection requirements and the system requirements as a whole so that adequate trade-offs can be made to ensure good system performance. This chapter discusses system-level requirements and basic difficulties that are associated with any high-speed obstacle detection system. Road geometry and how it affects the obstacle detection problem is also addressed. Finally, I examine how these affect the laser-based system in particular.

3.1 Safety Requirements

Kelly, in his Ph.D. thesis, codified the major safety requirements for a cross-country autonomous navigation system[24]. He described these requirements as falling into four categories: response, throughput, acuity, and fidelity. As his analysis showed, the performance limits for highway speed navigation are quite different from those for cross-country (relatively low-speed). Many of the assumptions that can be made in one scenario do not

hold for the other. However, the basic safety requirements are the same and we may use his formulation.

First, let us define a few terms. *Stopping distance* is the minimum distance needed to bring the vehicle to a full stop from the onset of braking. Although we may often be able to switch lanes upon detection of an obstacle in our current lane, our baseline mode of operation must give the vehicle the ability to come to a full stop before striking the object. Our first rule then states that the necessary *lookahead distance* for an obstacle detection system is the distance the vehicle travels in the time it takes to sense, process and detect an object, and apply the brakes, plus the stopping distance where v_0 is the initial velocity of the vehicle, a is the braking acceleration ($a < 0$) and t_{sense} , $t_{process}$, and t_{brake} are the latencies associated with sensor acquisition, computer processing, and braking reaction.

$$lookahead = v_0 t_{delay} + d_{stopping}$$

$$\Rightarrow lookahead = v_0 (t_{sense} + t_{process} + t_{brake}) - \frac{v_0^2}{2a}$$

Equation 3-1. To guarantee safety, the lookahead distance must allow for a vehicle at initial velocity v_0 to come to a full stop considering sensing, processing, and braking delays, and braking acceleration $a < 0$.

The second rule states that the throughput ratio, ρ_{cyc} , must not exceed unity (to guarantee coverage of the road), where the throughput ratio is given by:

$$\rho_{cyc} = \frac{v T_{cyc}}{\Delta R}$$

Equation 3-2. To guarantee safety, the obstacle detection system must be able to examine terrain at least as fast as it moves over it. This requires the throughput ratio ρ_{cyc} be smaller than 1 where T_{cyc} is the cycle time and ΔR is the projection of the pixels examined by the algorithm in a single cycle onto the ground plane (see Figure 3-1).

Equivalently, $1/\rho_{cyc}$ is the terrain oversampling factor. Oversampling factors greater than one may allow the system to track a potential obstacle over multiple frames for increased system reliability. ΔR is limited by the vertical field of view of the sensor,

although the algorithm may use only a portion of the available field-of-view. We may translate this throughput rule into a vertical field of view (VFOV) constraint (see Figure 3-2).




Figure 3-1. The vehicle must look ahead a distance L so that it can react in time. It must also examine enough terrain (ΔR) at each time step to guarantee coverage.

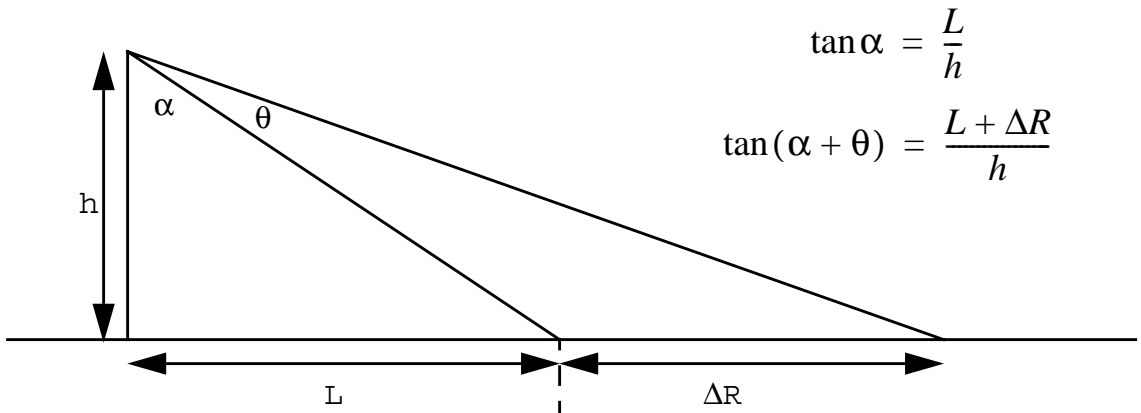


Figure 3-2. We can translate the throughput requirement on ΔR into a requirement on the vertical field-of-view (VFOV), shown here as θ .

In Figure 3-2, θ is the necessary VFOV for throughput requirements. We can solve the following equations for θ :

$$\tan(\alpha + \theta) = \frac{\tan \alpha + \tan \theta}{1 - \tan \alpha \tan \theta} = \frac{L + \Delta R}{h}$$

$$VFOV = \theta = \operatorname{atan} \frac{h \Delta R}{h^2 + L^2 + L \Delta R}$$

Equation 3-3. A minimum vertical field of view (VFOV) can be determined to guarantee sensor coverage of enough terrain.

Kelly uses two rules to determine the necessary acuity of the sensor. The first, which he calls the minimum acuity rule, is concerned with calculating the pitch of the vehicle when navigating rough terrain where there may be significant elevation differences

between the two vehicle axles. This will not be considered here since it is unimportant for highway applications where we plan to avoid anything that violates the flat ground plane assumption. The second states that at least two pixels must fall on a vertical obstacle for it to be detected in any single image. Throughout the thesis, we will consider $p = 20$ cm as the minimum obstacle height. The acuity rule can then be stated:

$$dz \leq \frac{p}{2}$$

where dz is the projection of a single pixel onto a vertical surface. We can relate dz to the vertical angular resolution of a sensor (see Figure 3-3) with the following equation:

$$\Delta\theta = \operatorname{atan}\frac{h}{L} - \operatorname{atan}\frac{h-dz}{L} \quad h, dz \ll L \Rightarrow \Delta\theta \approx \frac{dz}{L} \approx \frac{p}{2L}$$

Equation 3-4. The necessary vertical acuity of the sensor $\Delta\theta$ can be related to the minimum obstacle height p and the lookahead distance L .

In general, however, this greatly underestimates the necessary vertical angular resolution since oversampling is desirable for improved reliability of the detection system.




Figure 3-3. The vertical height a single measurement subtends, dz , is related to the vertical angular resolution of the sensor, $\Delta\theta$.

Since our scan speeds currently limit us to scans of only a single data line, we must instead detect a potential object with a single data line and then track the object over time. The object appears in a scan when it reaches the lookahead distance L . This relates the acuity requirement to the throughput requirement discussed previously. Given a sensor depression angle, an object of height p will be detectable in n successive frames where:

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

Equation 3-5. Assuming the obstacle is detectable in every line it is visible, an obstacle of height p will appear in n scanlines given a sensor scan rate of f , a sensor depression angle α , and vehicle velocity v .

Figure 3-4 illustrates the visibility limits on the obstacle. The shorter the obstacle, the fewer the data frames in which it will appear, and the less likely it will be detected. At highway speeds and a typical line scan rate of 40 Hz, a 20 cm obstacle will appear in approximately 16 frames before it disappears below the sensor field-of-view.




Figure 3-4. The laser first sees an object of height p at time 0. The laser moves up the obstacle surface as the vehicle moves towards the object until at time t_1 the laser passes over the top of the obstacle. The distance travelled before the object disappears depends on the laser inclination angle and the obstacle height p .

The last of Kelly's requirements, on system fidelity, concern the vehicle's ability to fit between two detected obstacles. We expect a very low density of obstacles on the highway (besides other vehicles), however, so we will ignore these rules here. In a highway system, we expect that the vehicle will either stop or move out of a lane if a static obstacle is detected in it.

It is instructive to calculate values for a typical highway scenario.

Table 3-1. Typical Values for Highway Obstacle Detection

Variable	Meaning	Value	Given (G) / Calculated (C)
v	vehicle velocity	26.7 m/s (60 mph)	G
a_{dec}	max. deceleration	-6.9 m/s ² (0.7 g)	G
t_{delay}	processing/braking delays	0.5 s	G
L	Lookahead distance	65 m	C
p	Obstacle height	0.2 m	G
T_{cyc}	Cycle time	0.3 s	G
ΔR	ground range	8.01 m	C

Table 3-1. Typical Values for Highway Obstacle Detection

Variable	Meaning	Value	Given (G) / Calculated (C)
h	height of sensor	1 m	G
VFOV	vertical field of view	0.0017 rad ($\sim 0.1^\circ$)	C
$\Delta\theta$	sensor vertical angular resolution	0.0015 rad ($\sim 0.1^\circ$)	C

It is interesting to note that, according to the calculated values in the table above, the necessary vertical field of view of the sensor need to be no more than the angular resolution of the sensor for the sample configuration. This is a rather surprising result and occurs because of the low grazing angles of the sensor sweep. This indicates that given a horizontal surface (except for obstacles), and a line scanner with a 0.1 degree spot size, we could guarantee coverage of the surface operating at only 3.3 Hz (although for reliability we probably want to oversample the area). Unfortunately, once we consider road surfaces with non-planar surfaces and vehicle pitch variations, we will see that it becomes more difficult to guarantee coverage with a single line-scan.

3.2 Road Geometry Considerations

Road geometry impacts an obstacle detection system in a number of ways. We first examine how horizontal road curvature affects the necessary horizontal field-of-view (HFOV).

3.2.1 Horizontal Road Curvature

The HFOV of the obstacle detection system must be, at a bare minimum, a lane width, w , at a distance L (the lookahead distance) on a straight road so that the vehicle may determine that its lane is clear. It may be argued for safety reasons that the system should also constantly monitor the adjacent lanes, or be able to see 3 lane widths, or $3w$, at a distance L . On a curved road, our HFOV may have to be even wider depending on whether we pan our sensor or not, as can be seen in Figure 3-5. A curve of 400 meter radius (typical of a

tight highway curve), can increase the necessary HFOV by a factor of 2 or more. Failing to steer the sensor would thus place even greater demands on the resolution of the sensor.




Figure 3-5. Unless the obstacle detection sensor is panned as a vehicle rounds a curve the horizontal field-of-view may have to be much wider. θ_s is the HFOV needed if the sensor is steered. θ is the minimum HFOV if the sensor is not steered. R is the horizontal radius of curvature, R . w is the width of the area of interest, usually either 1 or 3 lanes.

If the sensor is steered, the maximum HFOV occurs on a straight road, where:

$$HFOV = 2 \operatorname{atan} \frac{w}{2L}$$

Equation 3-6. For a sensor with a steered horizontal field of view (HFOV), the minimum HFOV that guarantees safety examines one lane-width w at the lookahead distance L .

To examine 3 lane-widths (12 meters), at a lookahead of 60 meters, the HFOV needed is 11.4° . If the sensor is not steered, however, the HFOV needed is affected significantly by the minimum radius of curvature of the highway, and can be described by:

$$HFOV = 2 \operatorname{atan} \left(\frac{2R + w - 2R \cos \beta}{2R \sin \beta} \right)$$

Equation 3-7. An unsteered sensor requires a larger HFOV to see a full lane-width w around a curve of radius R at a lookahead distance L .

Here β is determined by the lookahead and the radius of curvature. For the large R typical of highway curves, we can approximate $\beta = L/R$ where L is the lookahead (determined by the arc length of the curve). In this case, for a 60 meter lookahead, and a 300 meter minimum radius of curvature, the necessary HFOV is 22.7° . 300 meters is an approximate minimum radius of curvature based on highway design specifications for maximally-banked curves for a 60 m.p.h. design speed, or minimally-banked curves at a 50 m.p.h. design speed[1].

3.2.2 Vertical Road Curvature

Vertical curvature of the highway has a significant impact on small obstacle detection. One metric for highway design is the vertical crest *stopping sight distance*. The highway design manual allows a maximum road surface curvature (for a crest) by specifying that a driver in a vehicle must be able to see a 1/2 foot object in the road at the stopping sight distance L . The driver's eyes are assumed to be 3.50 feet above the road, and L is dependent on the designed speed of the road. For 60 m.p.h., the specified L is 580 feet. We can calculate the minimum radius of vertical curvature of the road using the following diagram and equations.




Figure 3-6. Highway design specifications state that a driver at height $h = 3.5$ feet above the road surface must be able to see an object of height $p = 0.5$ feet on the road at a distance L , called the stopping sight distance. L is dependent on the designed highway speed. Given L , we can calculate the radius of vertical curvature R .

From the diagram above, we can write the following equations:

$$x_1^2 + R^2 = (R + h)^2 \quad x_2^2 + R^2 = (R + p)^2 \quad x_1 + x_2 = L$$

Using the trivial assumptions $R \gg h$ and $R \gg p$, and solving the above equations for R , we have:

$$R = \frac{L^2}{2(h + p + 2\sqrt{hp})} = 25309 \text{ ft} = 7714.3 \text{ m}$$

Given the radius of curvature R , Figure 3-7 illustrates the difference between an assumed flat-world model and the actual curved road surface. Δh is the error in the elevation of the road surface at our supposed lookahead distance L (60 meters), and may be calculated:

$$\Delta h = R - \sqrt{R^2 - L^2} = 0.234 \text{ m}$$

Equation 3-8. A flat-world assumption will result in an error in the estimated road depth of Δh at our lookahead distance L for a radius of curvature R .

According to our calculation for Δh , we see that the deviations from flat world may be as large as the obstacles we are trying to detect in terms of elevation. Allowable changes in road grade for “sag” vertical curves may be nearly twice that for vertical crest curves [1]. This means that detection methods based on absolute elevation or deviation from a flat world model are inadequate.




Figure 3-7. Our sensor is normally inclined at angle θ to intersect a planar road at distance L . However, because the road has a finite radius of curvature, R , the ray from the sensor intersects at a distance L' , and at a depth of a below the assumed surface. Δh is the depth of the road below the assumed surface at the nominal lookahead distance L .

Although the flat world model may result in relatively large elevation errors (larger than some dangerous obstacles), a locally planar model, on the other hand, results in much smaller model errors. The best fit line of length L to a circular arc of radius R results in a maximum residual elevation error x , where:

$$x = \frac{L^2}{16R}$$

Equation 3-9. A best fit line of length L to a circular arc of radius R results in a maximum residual elevation error x .

So if we approximate a 40 meter section of a worst-case vertical crest curve road with a plane (which may be centered about the area of interest), we get a maximum model error of only 1.3 cm, which is much smaller than our obstacle size. Thus, a locally planar model appears to be adequate.

Vertical road curvature plus vehicle and sensor pitch variations can also make it very difficult to maintain a consistent lookahead distance without sensor field-of-view control. If we have a sensor aimed so that in the planar case it hits the road at our lookahead L , where does it hit the pavement if the road has a maximum curvature (with respect to highway specifications)? Based on Figure 3-7 we produce a somewhat simplified answer by first calculating a , the elevation error between the flat world and the point where our sensor ray actually intersects the road surface:

$$\tan \alpha = \frac{h+a}{L'} = \frac{h}{L}$$

$$(R-a)^2 + L'^2 = R^2 \Rightarrow a \approx \frac{L'^2}{2R} \quad \text{since} \quad a \ll R, a \ll L'$$

Substituting in for a , we have:

$$L' \approx \frac{h + L'^2/(2R)}{\tan \alpha} \Rightarrow L' \approx \frac{Rh}{L} - \sqrt{\left(\frac{Rh}{L}\right)^2 - 2Rh} = 95.74\text{m}$$

Equation 3-10. The actual lookahead distance L' may be significantly different from the supposed lookahead distance L (95.7 m as compared to 60 m) because of the road radius of curvature R .

This is a significant increase in our lookahead! Without sensor pitch or field-of-view control, a small variation in sensor pitch will also significantly change the lookahead distance (a pitch of 0.5° instead of 1° will double our lookahead distance), or can make the

sensor ray miss the road entirely. Pitch variations in the vehicle and sensor will likely force linear sensor-based methods to have cycling rates faster than the rate of pitch variation to guarantee coverage, depending on pitch variation amplitude.

3.3 Road Geometry Impact on Laser Reflectance

Laser reflectance provides us with no direct range estimates. Without range estimates, it can be difficult to determine what maneuvers are needed to avoid an obstacle, especially in the presence of other vehicles which may restrict our maneuvering. It is useful, therefore, to indirectly estimate target range.

To achieve fast data acquisition rates, we typically limit our laser sensor to providing only a single line of data at a constant depression angle. If our sensor is mounted at a 1 meter height, looking 60 meters away requires a depression angle of approximately 1 degree. It is easy to see that if vehicle pitch varies by more than this amount, and our sensor only provides us with a single row of data with a small IFOV, we may not see the road at all.

If we assume, however, that our sensor ray does intersect the ground and that we have a road model (given by local curvature) and sensor pitch estimate, then we could estimate the range of a target based on the first time we detect it as our vehicle moves. As the vehicle moves towards an obstacle, the sensor ray will first intersect the object at its base. The actual lookahead distance L' as illustrated in Figure 3-7 then provides us with our range estimate to the target when we first encounter it. Rewriting Equation 3-10 on page 38, with curvature $\kappa = 1/R$ ($\kappa > 0$) for vertical crest curves, we have:

$$L' \approx \frac{\tan \alpha - \sqrt{\tan^2 \alpha - 2\kappa h}}{\kappa}$$

Equation 3-11. The actual lookahead distance L' based on the road curvature κ , and the sensor depression angle α .

For sag curves as illustrated in Figure 3-8, $\kappa = -1/R$, we have the same equation where $\kappa < 0$. If we have a maximum sag curvature (double the magnitude of maximum vertical crest curvature), then L' is 44.56 meters when the lookahead is 60 meters.

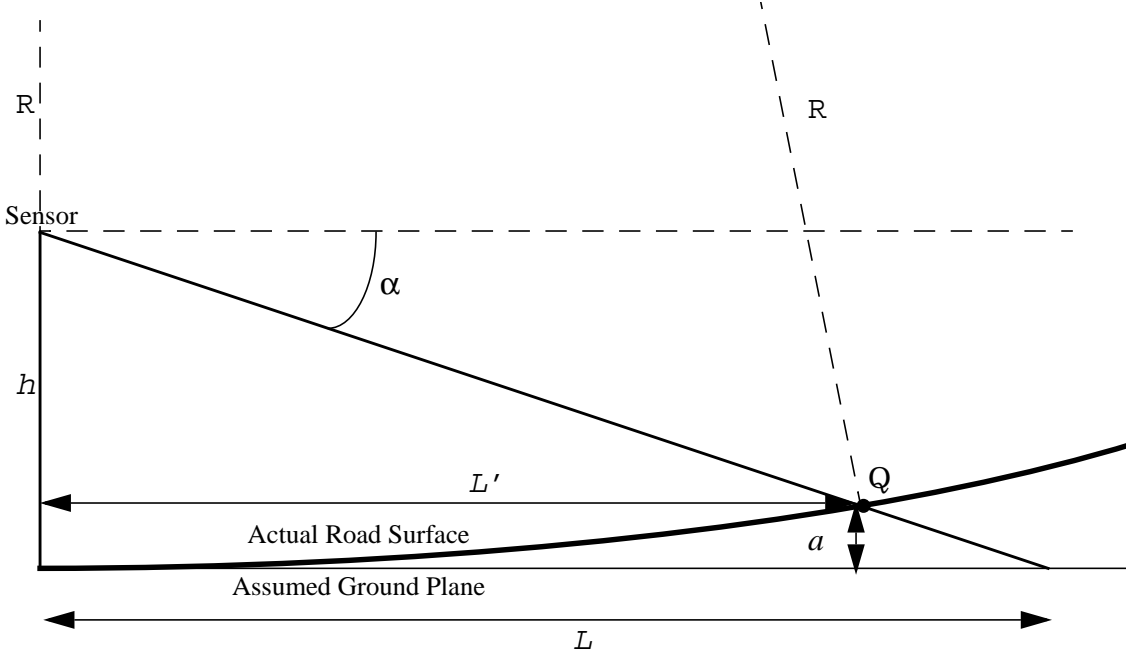


Figure 3-8. Geometry for a “sag” vertical road curvature. The sensor is mounted at height h , and looks at depression angle α and intersects the road at Q .

Given a range estimate when the target is first detected, the range estimate could be updated in subsequent frames by accounting for vehicle motion. Unfortunately, it is difficult to provide a decent estimate of curvature κ . Even if we have a very good estimate for α , the depression angle, we will have large errors in the range estimate described approximately by a Gaussian distribution with standard deviation:

$$\sigma_{L'} = \left(\frac{\partial}{\partial \alpha} L' \right) \sigma_{\alpha} + \left(\frac{\partial}{\partial \kappa} L' \right) \sigma_{\kappa}$$

Equation 3-12. Linearizing the function L' gives us a model for error in L' at a given α and κ , based on small gaussian errors in σ_{α} and σ_{κ} .

where σ_a and σ_k are the standard deviations associated with Gaussian errors in sensor depression angle and road curvature. Since the function for L' is highly nonlinear, the error is not truly Gaussian, and this only provides a very rough estimate for large errors. At the

nominal values of 60 meter lookahead and zero vertical curvature (a flat-world), then $\sigma_{L'}$ is approximately:

$$\sigma_{L'} = -3283\sigma_{\alpha} + 94017\sigma_{\kappa}$$

Equation 3-13. A localized model for the standard deviation of the lookahead distance error based on errors in sensor depression α and road curvature κ with small standard deviations σ_{α} and σ_{κ} . This linearization is based around $L = 60$ meters, $\alpha = 1$ degree, and $\kappa = 0$. Relatively small errors in either α or κ can result in large errors in lookahead distance.

Fortunately, even if our initial range estimate is poor, we can refine our range estimate to a target based on the reflectance information. A method to refine the range estimate is described further in Chapter 6.

3.4 Sensor Placement and Highway Design

For cross-country systems, sensors are placed as high as possible. The higher the sensor, the better it can see the area in front of it. In cross-country terrain sensing, occlusion is one of the most serious problems to conquer. The higher the sensor, the less the chance of occlusion: it is desirable to see over obstacles. Cross-country systems drive over smaller obstacles and drive around larger ones. They must also be able to spot negative obstacles, e.g. ditches or holes. To effectively spot such negative obstacles using range data, the sensor must be placed high so that it can look into the negative obstacle.

Because of the very long lookahead distances involved in highway driving, even very small obstacles occlude a significant amount of road. Fortunately, we expect a very low density of obstacles so we need not be concerned with occlusion in general. The problem of other vehicles occluding an obstacle will have to be dealt with, but this may be accomplished by allowing a vehicle to warn other vehicles behind it about upcoming obstacles. Our only requirement should be seeing the closest obstacle: the vehicle will choose to go around and not over obstacles. From the stereo equations presented in Chapter 2, we can see that the lower the sensor is placed, the more sensitive the method becomes. In fact, placing the sensor at obstacle height means that it must only determine whether it is imaging a point at a distance L or a point at infinite distance. Although placing the sensor low gives it a better view of vertical obstacles, it may render the same sensor unusable for road-fol-

lowing or other tasks since the sensor will have a worse view of the horizontal road surface. It is not unreasonable, however, to dedicate a sensor (or two) to obstacle detection. Placing the sensor too close to the ground will make it unable to cope with vertical curvature of the road -- the sensor won't be able to see over hills.

From the previous discussion on vertical road curvature, we see that a height of 3.5 feet (as outlined in the highway specifications) may be a reasonable minimum height for the sensor. Since an automated vehicle might have faster reaction time and automated vehicles might have better braking capability, the stopping sight distance, and thus the minimum height might be reduced somewhat.

Negative obstacles cannot be sensed at long distances with current range-based methods because of the shallow angles involved. Fortunately, negative obstacles generally grow gradually rather than suddenly. It may be reasonable to expect that an automated road can be well-enough maintained that there will be no seriously damaging negative obstacles such as potholes.

3.5 Summary

Stopping distance, throughput, and acuity constraints must be met to guarantee safety of an autonomous vehicle. Highway speeds demand large stopping distances which cause vehicle-mounted sensors to examine the road at low, grazing angles. Throughput requirements can be met fairly easily since a small sensor VFOV maps to a fairly large road area at low angles. System acuity must be moderately high, with a vertical resolution of approximately 0.1° or less, depending on the stopping distance.

Obstacle detection methods are highly sensitive to sensor and road geometry. Many previous systems have assumed a flat world model which does not seem justified for detecting small obstacles based on elevation. Locally planar methods, on the other hand, provide an adequate approximation for typical highway road curvatures.

Finally, it has been shown that small changes in road curvature or sensor pitch can drastically alter the current sensor lookahead distance. Either a large vertical field-of-view for

adequate oversampling or intelligent sensor pitch control to maintain lookahead distance will likely be needed for future highway-based obstacle detection systems.

CHAPTER 4 **Scanner Modeling**

This chapter will describe and characterize the scanner mechanism I used in my experiments and laser operational theory. It will also provide some performance metrics for the scanner system and discuss some of the problems associated with amplitude-modulated continuous wave laser devices.

4.1 Scanner Design

The scanner device consists of a two-axis mechanical scanning mechanism designed by K2T of Pittsburgh, PA, and a laser designed by Zoller and Fröhlich (Z+F) of Wangen, Germany. The scanner collects three dimensional information by steering laser energy through a 360 degree azimuth by 30 degree elevation field of view. The sensor produces a stream of 3D range points and intensity values arranged in a helical scan pattern.

The scanning device is comprised of a gold-coated aluminum mirror, a yoke assembly which allows the mirror to pivot vertically, a spindle assembly which rotates the yoke assembly horizontally, electronic circuits to establish the position of the mirror in space, electronic circuits to collect and store range data, a mechanical housing used to mount a

range detection sensor, and a commercially available range detection device (see Figure 4-2). The laser is located beneath the mirror assembly and points upwards. The scanner operates by continuously rotating the mirror in the two degrees of freedom. Data can be collected on the initial downward pivot of the mirror (downswing) and on the return sweep upwards (upswing). At each instant, the mirror deflects the laser beam at a known azimuth (determined by the rotation of the yoke assembly) and elevation (determined by the vertical position of the mirror). The laser travels outward and upon striking an object or reflective mass (such as dust or fog), some of the scattered laser energy returns to the detector along the same path. The range to the object is calculated by directly or indirectly measuring the time it takes for the laser signal to return.

One motor drives the yoke shaft which holds the mirror axle. The vertical position of the mirror is determined by a cam assembly. A cam-follower wheel rides on a variable elevation annular cam. The cam follower has a rack geared to the mirror axle, so that the local height of the cam surface determines the vertical mirror angle. The annular cam surrounds the yoke shaft and rotates, driven by a second motor. The difference in the rotation speeds of the yoke shaft and the cam surface determines the vertical nodding rate (see Figure 4-1).




Figure 4-1. The mirror deflection system. A rotating yoke shaft holds the mirror axle. Yoke rotation is provided by one motor and controls the azimuth direction of the laser beam. An annular cam with a variable elevation surface is spun by a second motor. A cam-follower wheel rolls on the cam surface. The cam follower has a rack which rotates the mirror axle through a gear as the follower moves up and down on the cam surface. For clarity, I have omitted the rigid connector between the yoke shaft and the cam follower assembly.

The maximum azimuth rotational speed of the scanner is 2400 rpm. Typical mirror nodding rates are 0.02 to 4.0 Hz. The frame rate is double the mirror nodding rate since a frame can be collected on each upswing and each downswing of the mirror. Depending upon the laser type used, data rates can reach 500,000 data points per second with an absolute accuracy of 1 to 2 cm with 2 to 8 mm resolution. The data can be formatted into either Cartesian or spherical coordinates in space as measured from the center of the scanning mechanism. Range can vary from 0.10m to 400m, depending on the type and power of the laser.




Figure 4-2. The laser scanner is on the right. A Sparc laptop, on the left, is used to interface with the scanner.

4.1.1 Laser Rangefinder Operational Theory

All laser rangefinders operate by emitting laser energy and measuring the time it takes for the energy to travel away from the sensor, strike a surface, and return. There are three basic methods for measuring this time interval: pulsed time-of-flight, AM phase-differencing, and FM beat frequency. For all methods, the target range is $z = (c\Delta t)/2$ where Δt is the roundtrip time of the laser energy, and c is the speed of light. Our scanning mechanism has been tested with several range detection devices, including two models of a pulsed time-of-flight laser spot sensor developed by Riegler and an Amplitude Modulated Continuous Wave (AMCW) laser range finder developed by Zoller and Fröhlich (Z+F) of Wangen, Germany. All three lasers operate in the near infrared wavelengths. A brief description of the operational theory for the three basic range detection methods follows.

4.1.1.1 Pulsed Time-of-Flight

A time-of-flight laser sensor operates by emitting a concentrated laser energy pulse at each pixel. The pulse travels away from the sensor, strikes a surface, and returns. A clock measures the time elapsed between the beginning of the pulse and the leading edge of the return pulse from the receiver (see Figure 4-3). Since lightspeed is 3×10^8 m/s, short distances can be measured only with very short pulses and complicated, fast electronics. The Riegl device we tested has an accuracy of ± 3 cm and a resolution of 2 cm with a maximum data rate of 12,000 points per second. Resolution and accuracy can be improved with longer dwell times by averaging many pulses and resulting in lower data rates. The maximum range of the laser is 150 meters for good reflectors ($\rho > 0.8$), or 50 meters for bad reflectors ($\rho > 0.1$). Retroreflectors can be detected up to 1000 meters away. The laser also provides 8-bit intensity for each pixel. The minimum distance the sensor can measure is 1 meter.




Figure 4-3. A pulsed time-of-flight laser sensor operates by emitting one or a series of closely spaced high-powered laser pulses and measuring the time until the laser pulse is detected by the receiver.

Another pulsed laser system by Riegl was created to see through dust clouds or fog. When operating in some environments, it is possible for some of the outgoing laser energy to be reflected by dust or fog while the remainder of the energy travels until it reaches a solid surface and returns to the detector. In this case, the receiver will see multiple return pulses. The laser measures the time elapsed between the beginning of the pulse and the *last* pulse returned to the receiver (see Figure 4-4). Provided the dust or fog is not too dense, this can find the range to the first solid surface. A last pulse detection method was imple-

mented in the laser system to create the high penetration system. This high-penetration version of the sensor has an accuracy of +/- 10 cm and a resolution of 10 cm and provides 8-bit reflectance data. Clearly more laser power is needed to detect targets through dust and fog at similar ranges since less energy reaches the final target.

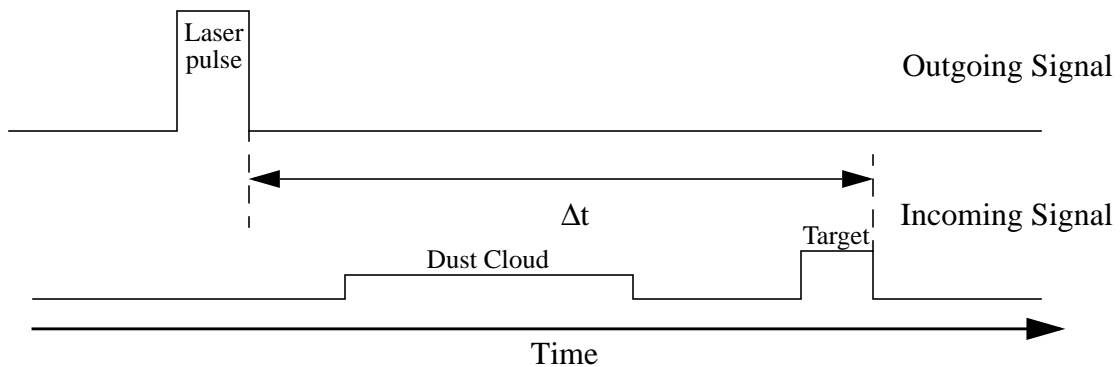


Figure 4-4. A high-penetration laser can be used to see through dust or fog. The sensor operates by transmitting a laser pulse and waiting for the last returned energy (up to some maximum time or distance) which is generally the solid surface of interest. Dust or fog may show up as some laser energy spread over time in the returning signal as the laser beam encounters small particles. Laser returns from dust or fog may even be larger in magnitude than those from the target.

4.1.1.2 Frequency-Modulated Continuous Wave (FMCW)

A second type of laser rangefinder is a frequency modulated continuous wave (FMCW) device. The emitted light is modulated by a sine wave at varying frequency, and mixed with the reflected energy. Range can be estimated by measuring the resulting beat frequency. Although other frequency modulations schemes are possible, the frequency “chirp” generally follows a periodic triangular waveform. With triangular frequency modulation, the distance to the target is proportional to the (maximum) beat frequency, i.e. the absolute difference in frequency between the returning signal and the emitted signal (see Figure 4-5).




Figure 4-5. Typical chirped waveforms and beat frequency for an FMCW laser. $f_e(t)$ is the emitted frequency, and $f_r(t)$ is the received frequency. B is the peak beat frequency for a measurement. Δt is the delay in the reflected signal and T_r is the period of the frequency chirp.

Range to the target can be calculated as follows:

$$z = \frac{c\Delta t}{2} = \frac{cBT_r}{4\Delta f}$$

Equation 4-1. Calculation of range z for an FMCW laser. c is the speed of light, B is the beat frequency, T_r is the period of the frequency sweep, and Δf is the change in swept frequency.

There are a number of competing factors in choosing system parameters for an FMCW laser. Note that the proper beat frequency B can only be measured during part of the time. The time of length Δt when it can not be measured we call the *dead time*. The fraction of dead time is $2\Delta t/T_r$, and the theoretical maximum measurable distance is $z_{max} = (cT_r)/4$ at which point B is measurable at only a single instant. To combat noise it is important to make T_r relatively large to keep the dead time as small as possible and increase the available sampling time to measure the beat frequency. However, a shorter T_r will increase the potential data rate of the sensor. In practice, the maximum range of the laser system is governed by the sensor power and noise levels, and does not come close to the theoretical maximum.

The minimum distance that can be measured by the system is dependent upon the amount of time it takes to properly sample the beat frequency signal. Bazin and Journet claim half of period of the beat signal is necessary, resulting in $z_{min} = c/(4\Delta f)$. Thus, a larger swept frequency will make it easier to detect small distances, but increasing the swept frequency will increase necessary filter bandwidths and can often introduce nonlinearities in the chirp waveform which, in turn, results in poor range estimates[3].

4.1.1.3 Amplitude Modulated Continuous Wave (AMCW)

The current CMU version of the laser scanner uses the Z+F AMCW laser. AMCW lasers operate by modulating the power of the emitted light with a sine wave of a given frequency. The returned energy waveform is delayed by the travel time Δt and appears proportionately phase-shifted when compared to the emitted energy. The distance is proportional to the phase, up to an ambiguity at a 2π phase difference:

$$z = \frac{\phi}{2\pi} \Delta R$$

Equation 4-2. The range, z , is proportional to the phase difference ϕ and the ambiguity interval, ΔR .

The ambiguity in phase results in a similar ambiguity in range of magnitude ΔR , called the ambiguity interval. A target at distance $x + \Delta R$ will be measured at range x .

The range and reflectance can be expressed in polar coordinates as:

$$I e^{i\phi} = \iint c I(\theta, \alpha) e^{i\phi(\theta, \alpha)} d\theta d\alpha$$

Equation 4-3. The intensity I and the phase ϕ of the returned signal are the result of an integral of signal energy over the instantaneous field of view (IFOV) of the laser. θ is the azimuth, and α is the depression angle.

If we assume that the phase is constant over the instantaneous field of view of the laser, then the resulting phase is the same as each component phase and the intensity is simply the sum of the individual contributions for each $d\theta$ and $d\alpha$. Undesirable results can occur

if this assumption is false. We discuss this possibility further in Section 4.2.1 on mixed pixels. Figure 4-6 shows the geometric relationship between the range and the intensity values.




Figure 4-6. Phasor diagram of received laser signal. I is the intensity, ϕ is the signal phase, and r is the target range. k is a proportionality constant.

The Z+F laser improves on the usual AMCW method by using a dual-frequency amplitude-modulated signal. The receiver measures the phase difference between the original and returned laser signal at both modulation frequencies (see Figure 4-7). The Z+F laser uses a combination of two modulation frequencies to provide improved accuracy over standard single modulation schemes. 12-bit accuracy in the phase shift measurement of the high frequency modulation (23 MHz) provides a range resolution of:

$$\Delta r = \frac{c}{2^{(b+1)}f} = 1.6 \text{ mm}$$

Equation 4-4. The range resolution Δr is determined by the speed of light c , the number of bits of accuracy in the phase differencing $b = 12$, and the modulation frequency $f = 23$ MHz.

The phase shift in the low frequency modulation (2.875 MHz) is used to disambiguate between multiple cycles in the high frequency, and allows for an overall ambiguity interval of:

$$\Delta R = \frac{c}{2f} = 52.2 \text{ meters}$$

Equation 4-5. The ambiguity interval ΔR is determined by the speed of light c and the modulation frequency $f = 2.875$ MHz.

The combination of the two channels effectively gives 15-bit range resolution over an unambiguous 52 meter range. The laser has an overall accuracy of ± 2 cm and can see up to 100 meters in the dark. At each pixel, in addition to the range measurement, the laser provides a 16-bit reflectance value which is determined from the high frequency channel. The Z+F sensor is very fast: it can operate at pixel rates as high as 500,000 per second. Using the Z+F laser, the scanner can provide images of up to 6000 columns by several hun-

dred rows over the 360 by 30 degree field of view with an angular resolution of 0.06° per pixel (see Figure 4-8). All of the data presented in this thesis uses the Z+F laser.




Figure 4-7. Range calculation is based on measuring the phase difference between the outgoing and the returned signal at two different modulation frequencies. The phase difference of the 2.875 MHz signal has an ambiguity interval of 52 meters and is used to disambiguate the 23 MHz signal which gives better range resolution.



Figure 4-8. Range and reflectance images taken in our lab with the scanner using the Z+F laser. We have subsampled each of the images in the horizontal direction by a factor of 3 to make them fit on to the page. In the range image, top, darker pixels are closer and brighter pixels are farther. The range values have been scaled for printing so that a range of 14 meters or more appears as white. In the reflectance image, bottom, brighter pixels correspond to points of higher reflectance. The reflectance values have been scaled for better printing as well.

4.1.2 Environmental Issues/Covering the Scanner

A feature of the Z+F laser is its *coaxial* design. Coaxial means that the laser transmitter and receiver share the same optical axis. A coaxial laser has the benefit of having the transmitter and receiver fields of view overlap at all distances which provides more efficient laser detection at all distances. Unfortunately, the coaxial design presented a challenge when it came time to seal the laser from the outside environment. The sensor needed a transparent cover to keep dust and dirt from collecting on optical and mechanical surfaces. Since the laser reflections are visible to the receiver at all distances, partial reflections from any cover interfere with the desired signal from a target. In contrast, if the receiver and transmitter fields of view only overlap at longer distances as in the ERIM laser scanner, objects at very close range (such as the cover) will be invisible to the receiver.

The first cover designed for the sensor was a cylindrical glass with an anti-reflective optical coating. The optical coating was designed to reduce laser reflections by 98%. We discovered, however, that the glass caused problems near the center of the vertical field-of-view. Because of the coaxial design, light reflected from the glass at normal incidences was directed back at the receiver. Although these reflections did not pose a problem for the high penetration pulsed laser system, the reflections did cause problems with the Z+F continuous wave system. At near-normal incidences (-5 to 5 degrees), the laser light reflected from the inside of the glass cover overpowered the signal returned from target surfaces resulting in poor range values and bright reflectance values towards the center of our laser images. We surmised that although the glass resulted in laser reflection at other incidences, these reflections did not present a problem since the reflected light was not directed back at the receiver aperture.

Our second cover was designed to avoid near-normal incidence problems. We opted for an uncoated, molded optical-grade plexiglass (type P-84) design in the shape of a truncated cone. The slope of the cone was chosen to avoid normal incidences over the entire vertical field-of-view of the laser. Since our vertical field of view could be adjusted to provide laser elevation from $+15^{\circ}$ to -25° , we chose the slope of the cone to be 30° . This shape avoided normal laser incidences over the entire field-of-view, with points above the horizon (the

least important areas for obstacle detection for terrain-navigating autonomous robots) having incidences closest to normal with a minimum incidence angle of 15° (see Figure 4-9).




Figure 4-9. The plexiglass cover was designed to protect optical and mechanical surfaces from dust without producing poor data from normal incidence reflections. This design provides a minimum incidence angle of 15 degrees anywhere inside the potential vertical field of view (VFOV).

Although the truncated cone design worked significantly better than the cylindrical cover, the cover still reflected some light back at the receiver which hampered laser performance. Another problem with the conical design is that it provides a variable incidence angle over the entire field-of-view. Since refraction and transmission losses through the cover are affected by incidence angle, these will also vary over the field-of-view. A better design would keep a constant incidence angle over the entire scanning range which directed reflections away from the receiver. An *equiangular spiral* ($r = ke^{a\theta}$ in polar coordinates) will generate such a surface. A new cover based on the spiral design is currently being investigated[31].

4.2 Performance

From a cursory visual inspection, the new scanner clearly provides better images than past scanners. However, it is instructive to reexamine some of the problems from past laser radar scanners and see how the new scanner fares. The first three subsections on mixed pixels, crosstalk, and range and reflectance precision are particular to the Z+F laser. The last subsection on angular precision is particular to the scanner mechanism itself.

4.2.1 Mixed Pixels

Previous AMCW laser systems had significant problems with mixed pixels. Mixed pixels are those that receive reflected energy from two or more surfaces. When the surfaces are separated by a significant distance, mixed pixels can result in reported ranges that are on neither surface, but somewhere between the two ranges, or even worse, either behind or in front of both surfaces. This is an inherent problem with AMCW laser radars and cannot be completely eliminated[19]. We illustrate the mixed pixel problem by considering the returned laser signal as a phasor as discussed previously, where the length of the vector corresponds to the signal amplitude or intensity, and the angle corresponds to the phase or range. If the laser receives energy from two surfaces, then the total energy received is just the sum of these signals which corresponds to vector addition in the complex plane. Based on the triangle inequality, the intensity of the total received signal is always less than or equal to the sum of the intensities from each surface.

In Figure 4-10 we show the case where the laser spot contains two surfaces which are separated by a phase difference of less than π (a range difference less than $\Delta R/2$). Summing the component signals from the two surfaces results in a signal with a phase, or range, between that of the phase (range) of the individual components.




Figure 4-10. A geometric interpretation of a mixed pixel for two surfaces separated by a phase difference of less than π . The overall signal is the sum of the components v_1 and v_2 . The resulting vector has a final phase (range) between the phase (range) of the two surfaces.

In Figure 4-11 we show a case where the laser spot contains two surfaces separated by a phase difference more than π (a range difference more than $\Delta R/2$). Depending on the rel-

ative strength of the component signals, the resulting signal can have a phase (range) either greater than or less than the phase (range) of both individual components.




Figure 4-11. In this case, the two signals are separated by a phase difference greater than π . This results in a final range less than that to either surface. It is also possible to get measure a range that is farther than either surface.

Fortunately, the Z+F laser reduces the frequency of mixed pixels by having a significantly smaller spot size than previous laser systems. The ERIM, for example, had a laser divergence, or instantaneous field-of-view (IFOV) of 0.5 degrees[10]. The Z+F, on the other hand, has a beam divergence of only 0.5 mrad (0.03°). Since the beam divergence is smaller than the scanner resolution, a pixel is less likely to fall on the edge of an object (see Figure 4-12) and thus cause an erroneous range value. When mixed pixels do occur, however, they are generally isolated pixels and may be removed by median filtering the range image. The higher resolutions provided by the scanner also make it easier to ignore or filter away isolated pixels without throwing out large amounts of the data.

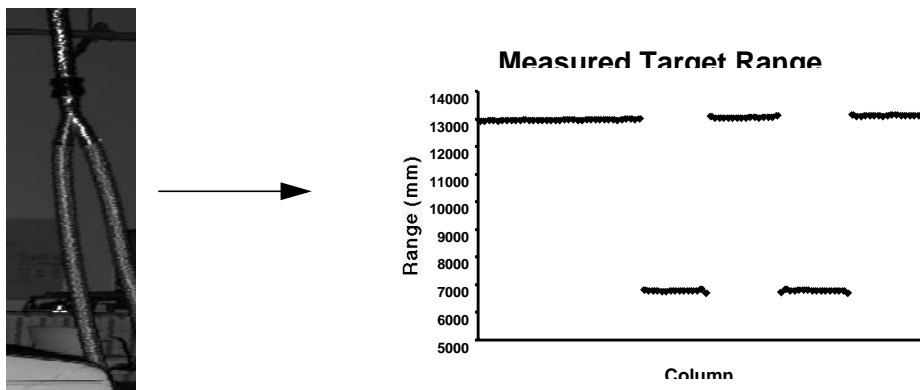


Figure 4-12. On the left is a portion of a reflectance image of an exhaust port. On the right is a graph of the corresponding range values of the indicated row in the reflectance image. Note that the range values immediately transition between the exhaust port and the far wall behind it. In this case, there are no mixed pixels resulting in erroneous ranges at the object edge.

4.2.2 Range/Reflectance Crosstalk

Ideally, a laser should measure the same median range to two targets which are equally distant but have different reflectance. In practice, however, this is not always the case. Crosstalk is a phenomenon in which reflectance or intensity values affect the measured range. To some degree, intensity always affects the range values. Lower intensity values decrease the signal-to-noise ratio and therefore increase the variance in measured range, although this should not affect the average range. As described in Hebert and Krotkov, the implementation of the electronics can play a factor in crosstalk[19]. Either a capacitive or inductive link between the intensity and range can cause such crosstalk.

The sensor optics can also produce effects similar to crosstalk caused by internal reflections. This is called optical crosstalk. I have found the plexiglass environmental cover causes significant optical crosstalk. Although internal reflections may be much reduced without the cover, they can still affect range and intensity measurements significantly when the target surface is far and the laser return from the target is of low energy. We can consider internal reflections (nearly constant) as a constant phasor added to the target signal. Figure 4-13 and Figure 4-14 illustrate the effect of the internal reflection.




Figure 4-13. A phasor diagram representing illustrating the errors caused by internal reflections. v_{ir} represents the portion of the signal returned from a constant internal reflection, and v_t represents the signal returned from the target. The measured return is v , the sum of v_t and v_{ir} . Here the measured intensity (the length of v) is greater than the intensity returned from the target.




Figure 4-14. A phasor diagram with the same internal reflection, v_{ir} . A target at a different distance is represented here by v_t . In this case the measured intensity of the resulting phasor v is less than the intensity of the returned from the target.

Of course, it is possible to compensate for a constant internal reflection electronically by subtracting the resulting phasor from new measurements. Often these reflections, however, are not constant. The light backscattered by the deflection mirror and the environmental cover depends on the angle of the mirror and the shape of the environmental cover. Although the phase of the internal reflection phasor may remain nearly constant (since the distance to the source of internal reflection is nearly constant), the magnitude of this phasor may change significantly. This makes perfect compensation for internal reflections nearly impossible.

Because of the modulation of the signal, internal reflections can either increase or decrease the measured intensity from the target, depending on its range. Ignoring signal effects other than internal reflection, the measured intensity of a target will oscillate with an interval equal to the laser ambiguity interval as it is moved along the optical axis. This effect is noticeable in Figure 4-15 where banding occurs in the intensity along the ground. While intensity should decrease monotonically as distance increases, the intensity on the ground seems to oscillate. This banding occurs at intervals of approximately 6.5 meters

which is the ambiguity interval of the high frequency modulation channel of the laser. This is caused by optical crosstalk.



Figure 4-15. A laser intensity image taken in a parking lot. Intensity should decrease monotonically as distance increases. Instead the intensity on the ground appears to oscillate as the distance increases. The banding along the ground is caused by optical crosstalk and occurs at approximately 6.5 meter intervals, the ambiguity interval of the high frequency modulation channel of the laser. At the left of the image, the white stripes with the sideways 'C' in-between them is a parking space ('C' stands for compact).

When targets are close and the returned signal is fairly strong, the Z+F sensor exhibits minimal crosstalk (optical or electronic), as shown in the graph in Figure 4-17 on page 62. Overlapping white and black posterboards were placed at approximately 2.6 meters from the sensor. The average range values reported for the black surface are approximately 2 cm greater than those for the white surface. Figure 4-17 also illustrates the problem of temporal mixing.

4.2.3 Temporal Mixing

Temporal mixing occurs when, within a single laser sample, the laser moves from one surface to another. This is especially noticeable at edges between darker and lighter surfaces. As a laser moves from a dark surface to a light surface, the returned signal amplitude

jumps. This can lead to large errors in phase estimation and thus large range errors (see Figure 4-16).




Figure 4-16. This illustrates the shape of the returned signal when an AMCW laser sweeps from a dark surface to a bright surface. At time t_0 , the laser transitions from a surface with 15% reflectivity to one with 100% reflectivity, resulting in a jump in signal amplitude. This frequently results in an incorrect phase difference measurement near the transition time, and can result in a drastically incorrect range measurement.

Notice that this phenomenon is not caused by the laser itself, and only appears when the laser is combined with a scanner. In the range and reflectance images below (see Figure 4-17) the effect of temporal mixing can clearly be seen at the edge boundary. Abnormally high range values are reported for quite a few pixels at edge boundaries. Some of these may be explained by low returned signal strength caused by specular reflections off of surfaces such as the metal struts on the garage door. However, the pixels on the edge between the black and white posterboard in the foreground cannot be explained by specular reflections, but are instead attributed to the temporal mixing of the laser signal. Temporal mixing can probably not be eliminated completely from any scanning AMCW laser system. Increasing

the modulation frequency and signal processing bandwidth or reducing scan speeds, however, may reduce the problem.

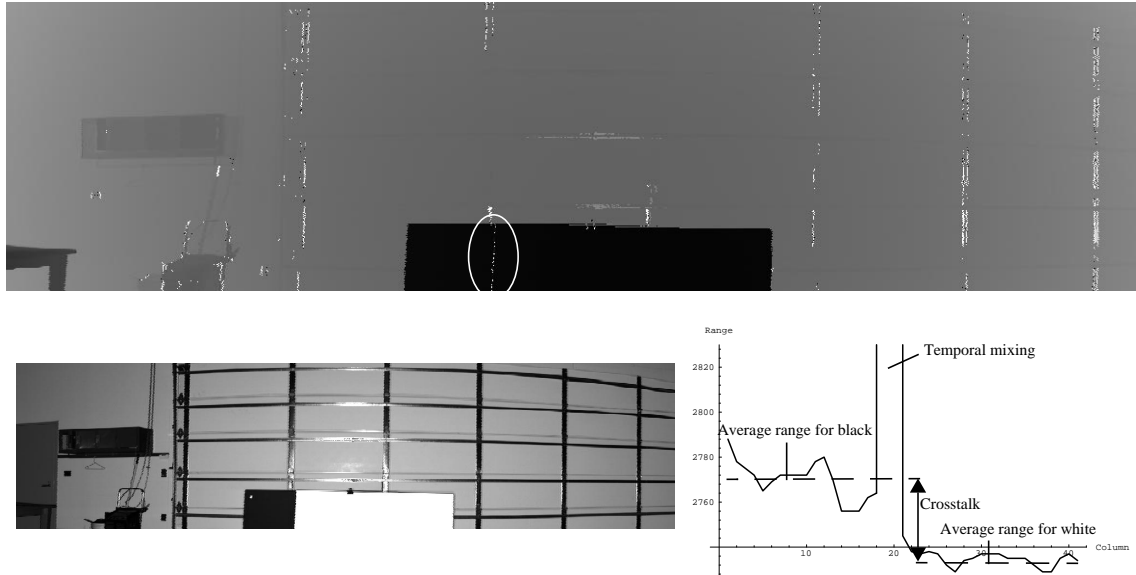


Figure 4-17. The range image, top, shows erroneous range values (the white pixels) corresponding to a number of edges in the reflectance image, bottom. In particular, notice the white pixels circled in the range image at the boundary between the black and white posterboard in the foreground of the reflectance image (bottom left). The posterboard pieces are placed at the same distance from the scanner. On the bottom right is a graph of the range values for a section of one row of the image centered on the black/white posterboard boundary. Notice the large momentary jump in range values at the border between the black and white surfaces caused by temporal mixing. Crosstalk results in a small (2 cm) difference reported range between the black and white surface.

4.2.4 Range and Reflectance Precision

To test the precision of the range and reflectance channels, data was taken with two different targets at increasing distances from the scanner. The scanning motors were turned off for this experiment. At each location a series of 900,000 measurements was taken from which the mean and standard deviation of the range and reflectance was recorded (see Table 4-1).

Table 4-1. Precision of Range and Reflectance

White Target				Black Target			
Mean Range (mm)	Mean Intensity	Range Std. Dev. (mm)	Intensity Std. Dev.	Mean Range (mm)	Mean Intensity	Range Std. Dev. (mm)	Intensity Std. Dev.
3603	31518	3.48	167.2	3622	3471	10.76	46.3
8100	13741	5.34	71.1	8087	1165	19.2	21.4
13783	5470	8.61	45.6	13737	537	29.9	14.5
23445	1958	14.78	27.7	23506	144	501.0	10.0
32398	1099	19.76	20.72	34715	122	3121	8.2
42103	600	28.9	16.3	47481	16.4	1115	6.0

Although the black target is of much lower reflectance (approximately only 10% as reflective), it provides similar range readings within several centimeters (within the possible errors of target placement) until beyond 20 meters. At that point, reflectance is low enough that the range values begin to degrade, as illustrated by the rapidly climbing standard deviation of the range measurement. Mean range values for the black target are within one standard deviation of the white target range value except for the last measurement.

4.2.5 Angular Precision

To measure the angular precision of the scanner, a set of 10 high-resolution laser images was taken at the same sensor location in a static environment. A set of 25 easily distinguishable points was chosen in the first image. The remainder of the images were then scanned by computer to find the corresponding points. Matching points were selected which minimized the sum of squared differences (SSD) in the reflectance for a small area in the matching images. The search window was centered at a point with identical azimuth and elevation. Since the precision was known to be reasonably good, the search window used was fairly small. The azimuth and elevation of each corresponding point was com-

pared with the azimuth and elevation of the original point in the initial image, and the error was recorded.

Because of the nature of the scan mechanism there is no guarantee that the points in each image will exactly overlap in elevation. Data can be acquired on both the downswing and upswing of the mirror. Rows were reversed for data collected on the upswing of the mirror so that areas of reflectance could be matched by the SSD method. To be sure that any errors reported were due to imprecision in the scanner mechanism and not the matching algorithm, some of the matching coordinates found by the algorithm were verified by hand. The computed coordinates were close (within a pixel) to the hand-picked matches.

A different approach for measuring angular precision is described by Hebert and Krotkov[19]. Their method found the centroid of multiple circular patches and compared the results over time. This method has the advantage that it did not rely on a matching algorithm to find corresponding points over time. However, since finding the centroid averages the errors in precision over an area, it may underestimate errors for single pixels caused by mechanical jitter.

For each feature in all images, the azimuth error had a standard deviation of 0.5 pixels (at maximum scanner resolution) or 0.03° . The elevation angle error had a much larger standard deviation of 0.3° , which corresponds to approximately 3.5 pixels at this resolution. Reported elevation angles for the same feature varies slightly more if the downswing elevations are compared to upswing elevation angles. This imprecision in the elevation has several potential causes. The first possible cause is backlash or slop in the elevation gear. A second cause is a manufacturing flaw in the cam surface which controls elevation. The bimodal cam is assumed to be perfectly symmetrical, but may not be. Explicit calibration for the two different sides might fix the problem and reduce elevation errors. A better solution would be to use a unimodal cam surface if this would not prove harmful to the balancing of the mechanism.

4.3 Summary

Laser scanners operate by steering laser energy over a scene and measuring target range at each angular position. Range measurement can be performed via by one of three methods: measuring the time-of-flight (TOF) of a discrete laser pulse; measuring the beat frequency of the mixed outgoing and return signals with a frequency modulated continuous wave system (FMCW); or measuring the phase difference between the outgoing and return signals with an amplitude-modulated continuous wave system (AMCW).

The laser built by Zoller and Frölich (Z+F) uses a dual frequency modulated AMCW method to provide a long ambiguity interval of 52 meters and a high resolution of 1.6 mm. Although it produces better data than any previous laser scanner at my disposal, there remain some problems. First, the coaxial design of the laser head made the design of an environmental cover difficult. Partial reflections from the cover are often seen by the receiver corrupting the desired return signal from the target. Second, AMCW methods present a number of difficulties inherent to the method. Erroneous range and intensity measurements can be caused by mixed pixels (where a pixel receives energy reflected from multiple surfaces), crosstalk (caused by internal reflections or other phenomena), and temporal mixing (where a sudden change in return amplitude may confuse the phase differencing calculation).

The scanning mechanism built by K2T uses a single mirror to deflect the laser beam in two directions. Angular precision tests have revealed good precision in azimuth deflections with a standard deviation of 0.5 pixels or 0.03° . Elevation precision is not as good with a standard deviation of 0.3° , or approximately 3.5 pixels at typical scanner resolution. Imprecision in the elevation may be caused by backlash or slop in the elevation gear, or a manufacturing flaw in the cam surface which controls elevation.

CHAPTER 5 **Reflectance Modeling**

Although the principle behind reflectance-based obstacle detection is simple, there are a number of difficulties with the method. Laser intensity is dependent on several factors: range, angle of incidence, surface albedo, and surface roughness. Exceptional values in any of these factors might lead to false detections or detection failures. Thus, it is important to have an accurate model of laser reflectance to improve the system performance and have a better understanding of where the techniques should succeed and where they will likely fail.

This chapter takes a primarily experimental approach to determining a better reflectance model. Thus, it focuses on modeling changes in reflectance based on external parameters and *controllable* internal parameters (such as integration time or sampling frequency). Although the data acquired with the sensor is dependent on a number of other factors such as detector responsivity, optical filter bandwidth, aperture size, etc., we will not attempt to build models that incorporate these parameters since this thesis does not focus on the laser design, but instead is concerned with the system design. In addition, most of these internal

parameters only affect the constant scale factors for the laser reflectance and the associated noise; these scale factors can be determined experimentally.

5.1 The Standard Reflectance Model

The standard laser reflectance model comes from laser communication theory. Assuming a Lambertian surface, the standard reflection model states[37]:

$$I \propto \frac{\rho \cos \theta}{z^2}$$

Equation 5-1. The standard reflectance model. I is the laser intensity, ρ is the surface albedo, θ is the angle of incidence, and z is the range to the target.

If this model were accurate for most objects, a vertical surface would provide a signal approximately 50 times stronger than the horizontal road surface at 60 meters due to the difference in the incidence angle (1° for the obstacle versus 89° for the road). In our experiments, however, we found vertical surfaces did not exhibit a fifty-fold increase in reflectance over horizontal surfaces. Instead, intensities were only a few times stronger for vertical surfaces. This discrepancy points to a problem with the angle of incidence term in Equation 5-1.

5.2 Verification of the Standard Model

Based on the preliminary data, it was clear that the laser reflectance was not adequately described by the baseline model for our purposes. I conducted a series of experiments to aid in better reflectance characterization. The first set of experiments sought to characterize noise in the laser signal. The second set of experiments was used to verify the inverse-squared range versus intensity relationship. The last set of experiments was used to better model the intensity versus incidence angle. This last experiment also helped to better characterize the noise.

5.2.1 Noise Experiments

There are two major types of noise that contribute to the signal. The first is *photon noise*, and is the result of ambient (non-laser) photons seen at the photodetector. Thus, it is

independent of the laser signal itself, and can be measured with an active photodetector and inactive laser. In daylight, photon noise can be significant. For our particular sensor, photon noise can overpower the signal received from some surfaces beyond 20 or 30 meters in strong daylight.

A second type of noise is what I term *signal noise*, since its amplitude is dependent on signal amplitude. This type of noise may be attributed to electronic noise in the laser modulation and/or noise in the acquisition and signal processing electronics.

I attempt to characterize both types of noise. Photon noise is independent of range and reflectance since it is independent of the actual signal. It may be reduced in comparison to the signal, however, by reducing the sampling frequency of the detector or increasing the laser signal strength.

5.2.1.1 Photon Noise

I performed several experiments to characterize the laser sensor with respect to photon noise. First, I characterized the ambient photon noise in indoor environments and outdoors environments at different times of day. The experiment was performed by turning off the laser but leaving the photodetector active and scanning the mirror over an environment. This gave an image of the photon noise in the scene. In every scene there were typically some bright spots in the noise image caused by light sources or specular surface reflections. We then calculated the mean μ_{PN} and standard deviation σ_{PN} of the intensity for each environment and recorded the maximum noise value detected both inclusive and exclusive of the unusually bright spots (such as light sources). I have tabulated the results below. Notice that the maximum may be more than one order of magnitude greater than the mean.

Table 5-1. Photon Noise Characteristics in Multiple Environments

Environment (all taken at $f = 250$ kHz)	Mean	Standard Deviation	Maximum (no light sources)	Maximum (incl. light sources)
Indoors (daytime, well-lit)	6.3	3.6	29	287
Outdoors, Sunny	72.0	37.9	296	344
Sky at Dusk	6.0	3.5	26	143

Table 5-1. Photon Noise Characteristics in Multiple Environments

Environment (all taken at $f = 250$ kHz)	Mean	Standard Deviation	Maximum (no light sources)	Maximum (incl. light sources)
Parking lot surface at dusk	3.4	1.9	14	143
Outdoors - Night	2.9	1.8	14	135

Next I characterized the photon noise with respect to the sensor sampling frequency. In this experiment, the sensor was aimed at a single point on a matte surface in an indoor or outdoor environment, and statistics were computed on approximately 300,000 samples taken at each frequency. Of course, the photon noise varies across a scene, depending on surface albedo and ambient lighting. As can be seen from the histograms below, the photon noise is Gaussian in nature (a single-tailed Gaussian), and the mean and variance increase with increasing sampling frequency (shorter dwell time). The increase in the mean of the photon noise does not indicate that more ambient photons are collected at higher sampling frequencies. Instead, the increase is a result of the decreased laser signal energy at higher sampling frequencies and thus a lower signal-to-noise ratio.

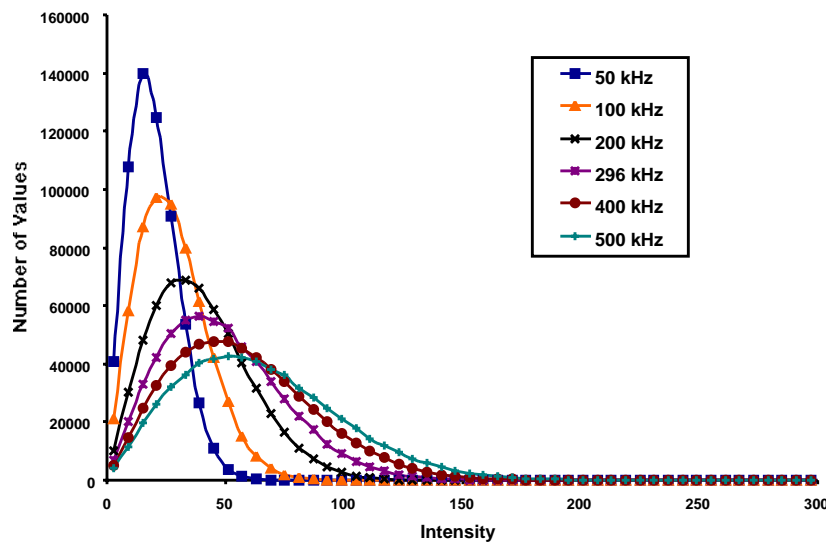
Photon Noise Histograms

Figure 5-1. Photon noise distributions for varying sampling frequencies. The data was sampled outdoors with bright sunlight on asphalt.

If we plot the mean-squared photon noise and the variance versus the sampling frequency (see Figure 5-2 on page 71), we can see that approximately, $\mu_{PN} \propto \sqrt{f}$ and $\sigma_{PN} \propto \sqrt{f}$. At distances over 30 meters, the photon noise can overwhelm the laser signal during bright daylight. For this reason, many of our experiments were performed indoors or at night. Table 5-2 characterizes the photon noise over 300,000 samples at various sampling frequencies for a typical indoor environment. For our typical sampling rates around 250 kHz, we can expect mean photon noise to be around 5, with maximum values under 30.

Table 5-2. Indoor Photon Noise versus Sampling Frequency

Frequency (kHz)	Mean	Standard Deviation	Maximum
50	2.2	1.3	10
100	3.2	1.8	15
200	4.6	2.6	25
296	5.6	3.1	27
400	6.6	3.7	32
450	6.6	3.7	32
500	7.2	4.1	35

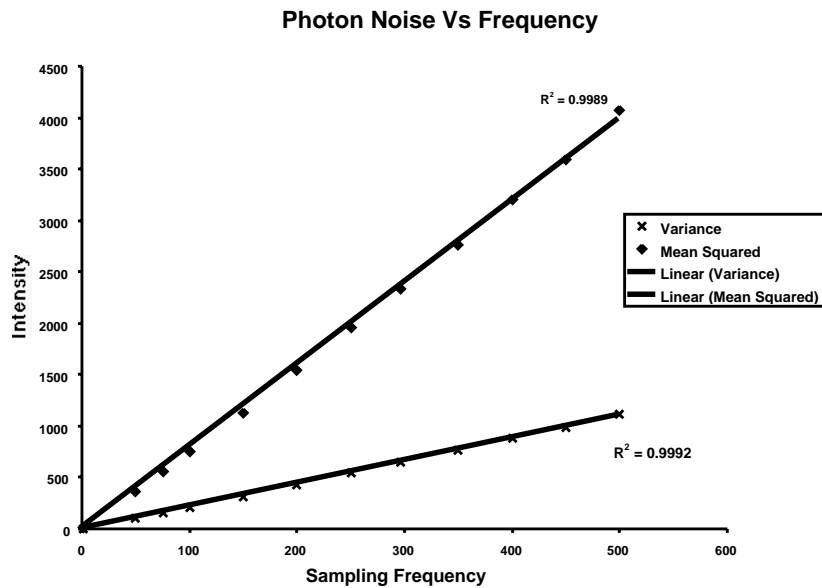


Figure 5-2. A plot of the mean-squared photon noise and the variance of the photon noise versus sampling frequency.

5.2.1.2 Signal Noise

Although it was relatively easy to avoid problems caused by photon noise corruption by performing many of my experiments at night or indoors, it was essential to have an accurate model of signal noise since signal noise was unavoidable. I first targeted a single point with the laser and collected a large number of data points to examine their distribution. Like photon noise, signal noise is approximately Gaussian. Throughout the rest of the thesis, I model signal noise as a Gaussian with zero mean and standard deviation σ_I , i.e. $G(0, \sigma_I)$. All experiments conducted to evaluate signal noise were performed indoors to minimize the effect of photon noise corruption. As the previous section indicated, photon noise is essentially negligible indoors.

There are two ways to improve the signal-to-noise ratio by increasing the signal energy. The first is to increase the signal power. The second is to increase the dwell time of the laser or decrease the sampling frequency. To illustrate the latter method, I collected approximately 600,000 laser intensity and range measurements of a single point at various sampling frequencies and computed the variance for each set of samples (see Figure 5-3 and

Figure 5-4). Note that the variance in both intensity and range is proportional to the sampling frequency.

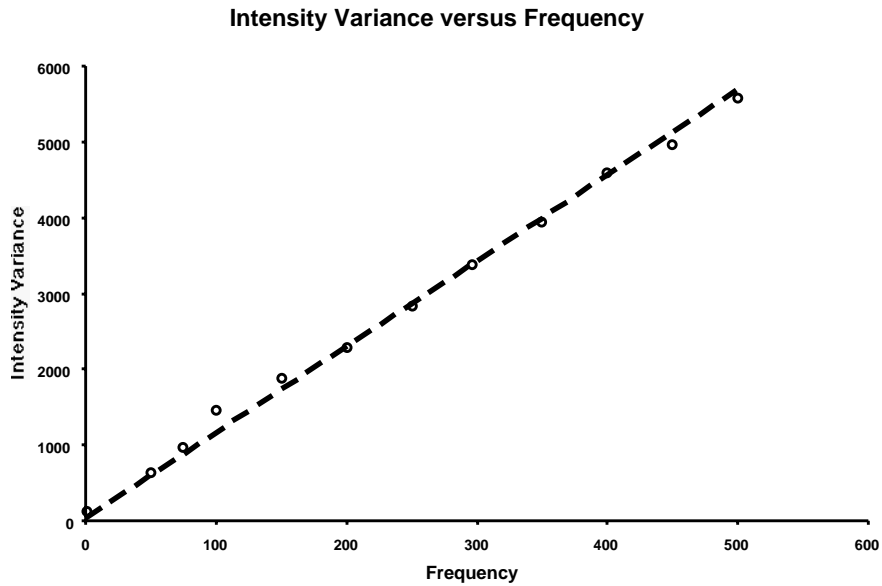


Figure 5-3. The effect of signal noise may be decreased by decreasing the sampling frequency. The variance in the measured intensity is proportional to the sampling frequency.

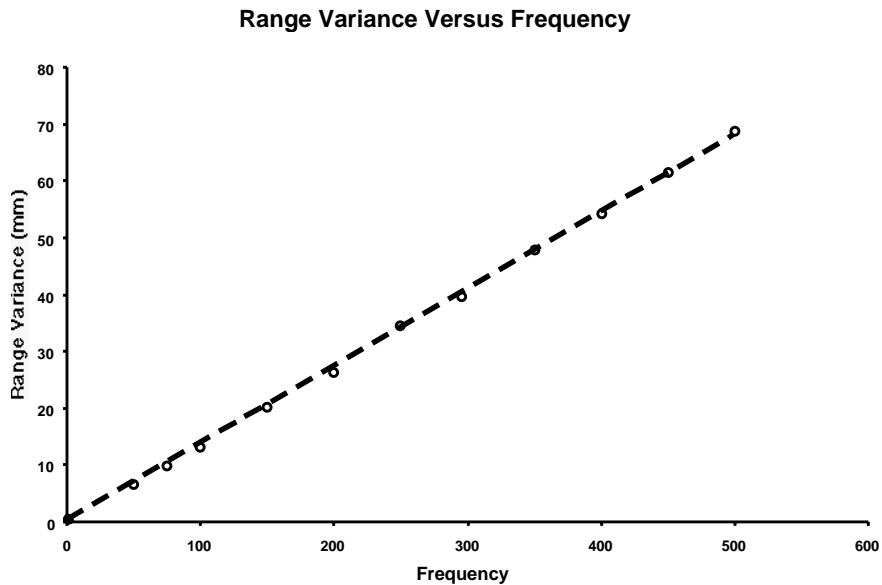


Figure 5-4. The variance in the measured range is proportional to the sampling frequency.

Of course, both methods for increasing the signal energy come at a cost. Increasing the laser power makes the laser less eye-safe, and decreasing the sampling rate decreases the data throughput of the sensor.

5.2.2 Reflectance versus range

As suggested by the standard model, reflectance depends heavily on range. It is important to understand and model this effect accurately to provide good detection and obstacle tracking capability at a variety of distances. To verify the model, I collected range and reflectance data from a non-scanning laser as I moved an object along the laser axis. The target was walked towards the sensor as it continuously grabbed data. I performed the test for multiple targets with similar qualitative results. I have graphed the data for a plywood sheet below. It should be noted that I used a 150 kHz sampling frequency.




Figure 5-5. This plot illustrates the relationship between range and intensity. The data was generated by moving a target (a plywood sheet) along the optical axis as the laser collected range and intensity information.

In the following figure, I show that the standard reflectance model provides a good description of the intensity versus range relationship. I multiply the intensity by the range

squared to give *normalized intensity*, and graph the result versus the range. Notice that it is nearly constant (albeit noisy) over most of the range.




Figure 5-6. The plot illustrates that normalized intensity (intensity times range-squared) is nearly constant beyond 7 meters. The normalized intensity is plotted on a logarithmic scale.

The normalized intensity drops at ranges less than approximately seven meters because of the laser optical design. Defocusing at the detector starts to occur at 7 meters. As the distance to the surface decreases, raw intensity continues to increase, but normalized intensity begins to drop. Defocusing becomes more serious as the object is moved closer, resulting in a drop in raw intensity at around 1 to 2 meters. With this caveat, I accept that intensity is inversely proportional to the square of the range. I do not attempt to model this defocusing effect, since I am concerned with detecting and tracking objects at much greater distances.

5.2.3 Reflectance versus incidence angle

From the initial reflectance data, it was fairly clear that for most of our objects, reflectance was not proportional to the cosine of the incidence angle as suggested by the standard reflectance model. This error may be attributed to two causes. First, there is a difference between the measurable macro geometry and the micro geometry of the surface. Rough surfaces provide reflectance values that do not correspond well to the macro geometry. Sec-

ond, some surfaces have specular reflectance qualities. To better quantify the effect incidence angle has on reflectance, I ran a set of experiments. I placed a variety of objects on a turntable and collected laser data as I rotated each object. At each angle, I collected 300,000 samples and calculated the mean and standard deviation of both intensity and range. As an alternative experiment, I could have scanned the laser across a flat surface and recorded reflectance and tracked the normalized intensity versus incidence angle, but that would have been subject to range errors and range noise. Graphs of the data appear in Figure 5-7 and Figure 5-12 through Figure 5-14 and are used to illustrate the reflectance models described in Section 5.2.4.

It should be noted that there are at least two significant sources of error in the data. The first is caused by surface albedo variations. Since I could not place the surface exactly at the center of rotation of the turntable, the laser moved along the surface as it was rotated. So at each different angle, the laser was sampling a slightly different surface point. Thus, local variations in albedo show up as noise in the graphs. The second source of error was the difficulty of measuring the absolute angle of incidence. The absolute angle as illustrated in the graphs is only accurate to within 2 degrees. Relative angles, however, are accurate to within 0.1 degrees.

The first two graphs show the variations in reflectance for two smooth surfaces. The first, a sheet of matte white posterboard matches the Lambertian cosine curve quite well. The second matches a Lambertian curve at high incidence angles, but there is clearly some specular reflection in the center. Since the Lambertian model was inadequate, I investigated other reflectance models.

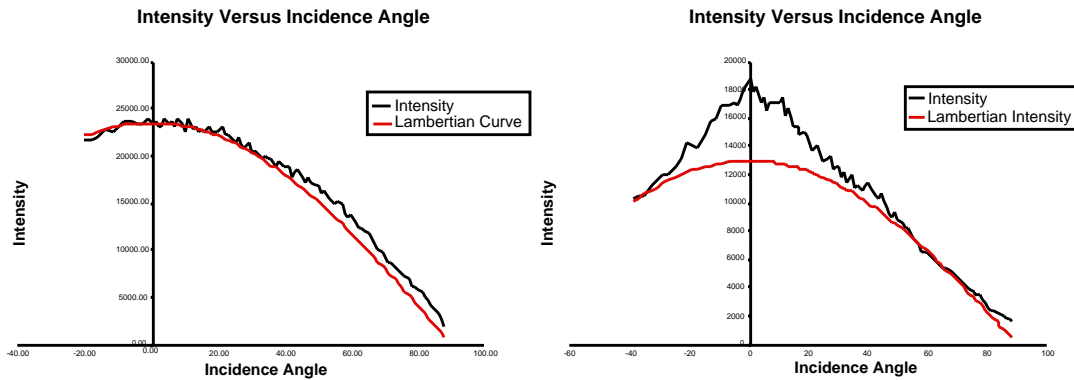


Figure 5-7. Intensity versus incidence angle curves for two smooth surfaces. In each case, I have fit a Lambertian cosine curve to the data to illustrate the fit. The graph on the left shows data for a piece of white posterboard. The graph on the right is a cardboard box. Note that the cardboard surface is somewhat specular; roughly a third of the intensity at normal incidence is the result of specular reflection.

5.2.4 Reflectance Models

There are a number of reflectance models to be found throughout the computer graphics and computer vision literature. Besides the Lambertian model, I look at two models in particular. The Torrance and Sparrow model is frequently used to model specular reflections from rough surfaces. More recently, Oren and Nayar developed a reflectance model that works for many rough, diffuse surfaces. Both models are bidirectional reflectance functions, i.e. they model reflectance with regards to two directions, the direction of incidence and the direction of reflection. Fortunately, in laser-based vision, the incident and reflected rays are coincident and we may greatly simplify both reflectance models to unidirectional

models. See Figure 5-8 and Figure 5-9 for a comparison of generic vision and laser-based vision reflectance diagrams.




Figure 5-8. A bidirectional reflectance diagram for a typical light source-camera setup. The incident beam direction is specified by its zenith angle θ measured from the surface normal. The direction of the reflected flux is specified by two angles, the zenith angle ψ measured from the surface normal, and the azimuthal angle ϕ measured from the plane of incidence.




Figure 5-9. A unidirectional reflectance diagram for laser-based vision. The direction of the incident beam and reflected flux can be described by a single angle, the zenith angle measured from the surface normal.

It is important to realize that the distance or magnification will affect the reflectance from a rough surface. The closer we are to the surface, the rougher it appears (unless truly fractal). Both the Torrance and Sparrow and Oren and Nayar models assume that the surface is made of long, symmetrical V-shaped cavities. “Long” means that the cavity length is much greater than the width. Each model is based on geometrical optics, and is only valid

if the wavelength is significantly smaller than the surface facets. In addition, it is assumed that the facet area is small compared to the area imaged by a single pixel (see Figure 5-10). For our purpose, the facet area must be smaller than the laser spot size.




Figure 5-10. The surface model assumed by the Torrance & Sparrow and Oren & Nayar reflectance models. n is the macro surface normal, a is the facet normal, da is the facet area, and dA is the area imaged by a pixel or laser spot.

The above assumptions are summarized as:

$$\lambda^2 \ll da \ll dA$$

For each of the models described below, we show some reflectance versus incidence angle data taken for a real surface. Ideally, I would have gathered this data at a range similar to that of the potential obstacles, i.e. approximately 60 meters. Unfortunately, at 60 meters, the signal-to-noise ratio is much lower and would have resulted in poor experimental data. Instead, I took most of the data at close ranges (around 2 to 4 meters), to maximize the returned signal strength. For the surfaces which fit either model or the combination, the spot size was significantly larger than the facet size, so these surfaces should behave similarly at 60 meters distance.

5.2.4.1 Torrance and Sparrow Reflectance Model

I have simplified the Torrance and Sparrow model from its more general bidirectional form for the coincident laser case. In this simplified form, the Torrance and Sparrow model

appears similar to the Lambertian model, adding a specular lobe to the Lambertian model at incidence angles near zero:

$$I = k\rho \left(C \cos \theta + \frac{1-C}{\cos \theta} e^{\frac{-\theta^2}{2\sigma^2}} \right), \quad 0 \leq \theta \leq \frac{\pi}{4}$$

$$I = k\rho \cos \theta \left(C + 2(1-C)e^{\frac{-\theta^2}{2\sigma^2}} \right), \quad \frac{\pi}{4} < \theta \leq \frac{\pi}{2}$$

Equation 5-2. The Torrance and Sparrow reflectance model. The model adds a term to the Lambertian model to account for specular reflections near zero incidence angle. Here σ is a measure of surface roughness and C is the percentage of energy reflected diffusely.

Notice that outside the specular lobe, the model converges to the Lambertian model. The Torrance and Sparrow model does a good job of modeling fairly smooth specular surfaces such as our cardboard surface (see Figure 5-11).

While some smooth surfaces can be modeled fairly well by the Lambertian or Torrance and Sparrow models, reflectance at high incidence angles from rough surfaces can not be adequately described by either model. Real surface intensities do not fall off as rapidly as these models predict with increasing incidence angle.




Figure 5-11. The same intensity versus incidence angle graph for the cardboard in Figure 5-7 on page 77. I have added a Torrance and Sparrow model curve to this graph. The Torrance and Sparrow model achieves a much better fit than the Lambertian curve by accounting for the wide specular lobe.

5.2.4.2 Oren and Nayar Reflectance Model

Oren and Nayar illustrated that a rough surface with Lambertian facets will result in non-Lambertian reflectance. The Oren and Nayar model assumes a surface made of small Lambertian planar facets with varying orientations according to a known distribution. “Small facets” in this case means much smaller than a pixel, or in our case, smaller than the laser spot. Note that this depends on the range to the target: at close ranges, the surface will appear rougher than at longer ranges. The Oren and Nayar model in its simplified, coincident form can be stated:

$$I(\theta) = k\rho(\cos\theta + \varepsilon\sin^2\theta)$$

Equation 5-3. The simplified Oren and Nayar reflectance model. k is a constant, ρ is the albedo, θ is the angle of incidence, and ε is a measure of surface roughness.

At angles near the normal, the Oren and Nayar model approaches the Lambertian model, but it increases the reflectance at high incidence angles. The rapidity of the drop-off in the reflectance at high incidence angles is controlled by a surface roughness parameter, ε . This model is a good approximation for many non-specular rough surfaces, such as an unfinished cement-like surface shown below.




Figure 5-12. Reflectance versus incidence angle for a concrete patch surface. The Oren and Nayar model provides a good approximation to the data with a surface roughness parameter $\varepsilon = 0.4$. A Lambertian model incorrectly predicts a more drastic drop in reflectance with increasing incidence angle.

5.2.4.3 Combining the models

The Torrance and Sparrow and Oren and Nayar models use an identical model for surface structure. The difference between the two models lies in the assumptions about individual facet reflection. Since the Torrance and Sparrow and Oren and Nayar models improve the reflectance modeling over different incidence ranges, it is straightforward to combine them to take advantage of both.

$$f(\theta) = C(\cos\theta + \varepsilon \sin^2\theta) + \frac{(1-C)}{\cos\theta} e^{\frac{-\theta^2}{2\sigma_\alpha^2}}, \text{ for } 0 \leq \theta \leq \frac{\pi}{4}$$

$$f(\theta) = C(\cos\theta + \varepsilon \sin^2\theta) + 2(1-C)\cos\theta e^{\frac{-\theta^2}{2\sigma_\alpha^2}}, \text{ for } \frac{\pi}{4} \leq \theta \leq \frac{\pi}{2}$$

Equation 5-4. A reflectance model that uses the Torrance and Sparrow specular component and the Oren and Nayar diffuse component to account for changes in reflection based on incidence angle θ . C is the percentage of energy reflected diffusely from a surface. ε is a surface roughness parameter, and σ_α is the standard deviation of facet slope.

The combined model works for many surfaces, such as a sheet of plywood (see Figure 5-13). The Torrance and Sparrow model provides a good fit at low incidence angles, and the Oren and Nayar model provides a good fit at high incidence angles. Combining them provides a good fit over the entire range.




Figure 5-13. Intensity versus incidence angle for a sheet of plywood. The combination of the two models provides a good fit over the entire incidence angle range.

There are some surfaces for which the data (at the current magnification factor) do not fit any of our models well. One such surface is asphalt. The facets on the asphalt are larger than the laser spot size at close ranges (2 to 4 meters). A sample reflectance curve for an asphalt surface is given in Figure 5-14. It should be noted that this curve is largely random. Since the laser spot is smaller than a facet, pointing the laser at a different point on the surface will likely result in a very different set of data.




Figure 5-14. Intensity versus incidence angle for a small area of asphalt at a range of 1.5 meters. Given the large facets in asphalt, the intensity is largely random on a global scale.

Although none of our reflectance models can simulate the asphalt at this scale and such a close range, they can provide a reasonable fit at longer ranges (at 50 meters) where the spot size becomes significantly larger than the facet size. The weakness of the signal at long ranges, however, makes it difficult to generate such a curve. However, we can simulate the intensity versus incidence angle relationship at longer ranges by using a larger laser spot size or averaging the reflectance over many pixels on the surface. I placed an asphalt sample approximately 50 cm by 50 cm in size at a range of 13 meters, and collected images as I rotated it through 90 degrees. I then calculated the mean intensity of all the pixels that fell on the asphalt sample for each image. This generated the plot found in Figure 5-15.

Note that this data is subject to significant noise especially at the grazing angles since very few pixels fell on the surface at this range.




Figure 5-15. Intensity versus angle for an asphalt sample using intensity averaged over the entire sample.

Not surprisingly, another surface which is not modeled well by our reflectance models is a piece of deer hide (see Figure 5-16). Intensity is quite noisy, and the maximum signal is less than three times stronger than the weakest signal. The maximum reflectance does not occur near the normal, but instead near an incidence of 40 degrees. Again, given the large scale surface roughness, I expect that a different spot on the hide would generate a significantly different curve.




Figure 5-16. Intensity versus incidence angle for a deer hide. Intensity is fairly noisy, but the maximum to minimum intensity ratio is small.

5.2.5 Signal Noise Implications

The intensity versus incidence angle experiments also provide us information about the signal noise by examining the signal variance at each point. If we plot intensity variance versus intensity (see Figure 5-17), we find that we can model the signal noise variance as being roughly proportional to the mean intensity. Data taken with varying range and surface albedo at a single sampling frequency also supports this model as a good approximation for the signal noise variance. Though the relationship may not be exact, it does allow us to make some predictions about signal noise at different signal levels. As noted earlier, signal noise is also proportional to the sampling frequency, so I model signal variance (since mean intensity is independent of the sampling frequency):

$$\sigma_I^2 \propto \mu_I f$$

Equation 5-5. Signal variance is proportional to the product of the mean and the sampling frequency.




Figure 5-17. Plot of intensity variance versus mean intensity as the white posterboard surface was rotated through 90 degrees. The intensity variance is proportional to the intensity.

5.3 An Improved Reflectance Model

The previous experiments provide a more complete reflectance model for our laser sensor:

$$I(\theta) = \frac{k\rho f(\theta)}{z^2} + G(0, \sigma_I) + G(\mu_P, \sigma_P)$$

$$\sigma_I \propto \sqrt{\mu_I f}$$

Equation 5-6. A more complete reflectance model based on our experiments. I have changed the intensity dependence on the incidence angle, and have explicitly accounted for both signal noise and photon noise, modeled as gaussian distributions. $G(\mu, \sigma)$ represents a gaussian distribution with mean μ and standard deviation σ .

As noted in Section 5.2.2, this reflectance model is only accurate for ranges greater than approximately 7 meters for our particular sensor.

In the next chapter, I use this reflectance model to help detect obstacles in mild terrain such as parking lots or highways. I also use the range dependency in the intensity to provide range estimates without using the range data from the scanner. First, there is one last important physical effect on the laser intensity to consider: signal drift.

5.4 Intensity Drift

Signal drift can be a problem in almost all sensor systems. Laser systems are no different. Drift in laser systems is primarily caused by changes in temperature. Temperature changes can affect the laser system in a number of ways. The main effect, however, is that changes in temperature change the length of the laser cavity and thus the laser frequency. For laser ranging systems, the change in laser frequency can show up as a shift in phase resulting in a range drift. For intensity measurements, changes in the laser frequency affect the gain of the optical filter in front of the receiver. As the frequency shifts farther from the center of the optical bandwidth filter, the measured intensity of the signal will drop.

I observed laser intensity drift by pointing the laser at a fixed point and measuring the intensity over time. Every 15 seconds, 30,000 intensity measurements were taken and averaged. A plot of intensity over time is given in Figure 5-18. During this time, the temperature inside the scanner rose from approximately 20°C to approximately 40°C. I did not have a direct means of recording the diode temperature.




Figure 5-18. Intensity drift over time. Intensity drops significantly (approximately 1/3) over approximately 2 hours as the diode temperature rises.

5.5 Chapter Summary/Conclusion

Laser reflectance depends on many factors including: target range, angle of incidence, albedo, surface roughness, and shininess (how much energy the surface reflects specularly). While it is unlikely that all of these parameters can be determined from a limited number of reflectance measurements from an unknown obstacle, the model can be simplified for the purpose of target range estimation. The next chapter will show how target range can be estimated from multiple intensity measurements over time.

While the full reflectance model cannot be fully exploited in the detection of obstacles, it is useful to understanding where intensity-based segmentation methods will fail. Clearly, there is a set of surfaces that when vertical will appear nearly identical to the horizontal road. Examining the models may show which surfaces can “spooof” the road. For example, an obstacle that has a significantly lower albedo than the road may look very similar to the road because of the difference in incidence angle. It is important that these same surfaces can be detected by a different obstacle detection method. Alternatively, these same surfaces might be detectable with a different laser wavelength since the albedo is wavelength-dependent.

CHAPTER 6

Obstacle Detection and Tracking

The experiments detailed in Chapters 4 and 5 provide us with a good understanding of both scanner and laser. Now I examine laser intensity-based obstacle detection methods in more detail. As stated before, for a vehicle-mounted laser, vertical obstacles should appear brighter than the horizontal road at long distances on the basis of the significantly different incidence angles. In this context, intensity-based obstacle detection is essentially a search for unusually bright spots in reflectance images.

6.1 System Overview

The detection system designed consists of a vehicle-mounted scanning laser, a vehicle position sensor, and set of software algorithms which use the laser intensity to perform obstacle detection and tracking.

6.1.1 Hardware

The laser scanner uses a 2-DOF (degrees of freedom) mechanism designed by K2T and a laser designed by Zoller + Fröhlich, and is more fully described in Chapter 4. Although the scanner mechanism is 2-DOF, the detection system only uses 1 DOF. The scanner vertical field of view (VFOV) is 30° which is far too large for highway-based detection. Since the scanner provides no method to scan only part of this VFOV or computer-control the vertical position of the mirror, I chose to turn off the vertical scanning completely. The result is a scanner that provides a 360 degree single line laser scan at 40 Hz. Only a small amount of the 360° scanline (approximately 60° in front of the vehicle) is kept and processed.

The position sensor on the vehicle (Datron DLS-1) uses optical correlation to provide vehicle motion estimates along the longitudinal axis of the vehicle with 0.1% accuracy.

6.1.2 Software

The software uses the returned intensity of the laser to classify pixels as obstacle or non-obstacle in each line of data. Since any detection method is likely to generate some false positives, the system attempts to verify all detections through temporal filtering and tracking. I assume that all detected obstacles are static (I am only concerned with static obstacles in this thesis). In a commercial system, any detected obstacles would need to be verified with a vehicle detection system (such as radar) so that a non-static obstacle could be identified as such.

Assuming an obstacle is static, it is simple to predict its azimuth in one frame given its azimuth and range in the previous frame (or the last frame in which it was detected), and a model of the egomotion of the vehicle. By restricting the system to use only the laser intensity (in the hopes of making a cost-effective system), it does not have direct range information. If we assume that we detect the obstacle when the laser first intersects it at its base, we can estimate the initial range to the object based on the inclination angle of the laser beam. Although we know the sensor-relative laser pitch, our current setup does not provide us

with vehicle pitch relative to the ground or vertical curvature of the road, so our range estimate is very rough.

While segmentation provides an accurate picture of where obstacles are in azimuth, it would be useful to both the tracking system and a vehicle control system to know the range of each of the obstacles. While range estimates based on the sensor inclination angle may be useful, they are fairly unreliable. I have developed an alternative means of estimating the range from the intensity change over time as the vehicle approaches the obstacle.

The current intensity-based obstacle detection system consists of 4 modules (update, classification, merge, and range estimation) and 2 lists. The “active” list contains all potential obstacles that have been detected or confirmed recently, and the “inactive” list contains all potential obstacles in front of the vehicle which have not been seen recently, possibly because they are out of the VFOV of the sensor. A diagram of system architecture is given below, and each module is described in more detail in the following sections.

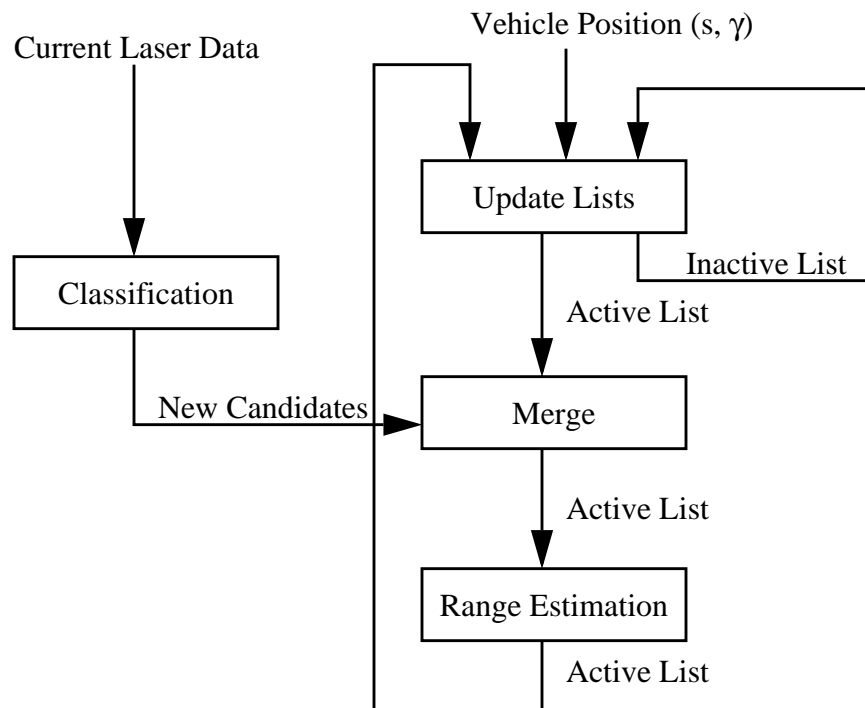


Figure 6-1. The detection system architecture consists of 4 modules (update, classification, merge, and range estimation). These modules maintain 2 obstacles lists, labelled “active” and “inactive.”

6.2 Update

At each step, the laser collects one line of data, and a position sensor records the change in vehicle position and yaw since the previous scan as a vector $Y = (s, \gamma)$. Obstacles on the active list that have not been confirmed recently are moved to the inactive list. Objects on the inactive list that are no longer a threat (the vehicle has passed them) are removed from the list.

The update module uses the measured change in vehicle position to update the expected vehicle-relative positions of all remaining obstacles on the active and inactive lists as shown in the figure below. Obstacle state parameters are given by a vector $X = (R, \phi)$.

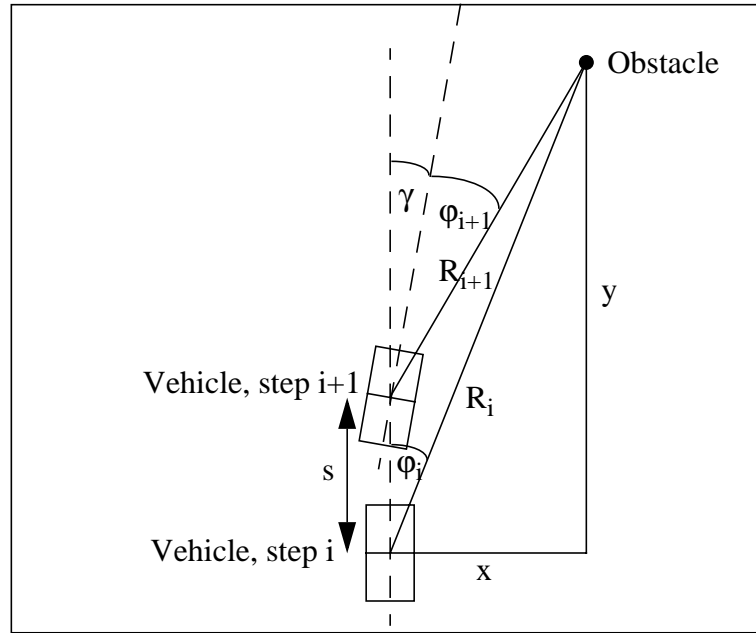


Figure 6-2. The system updates the estimate of an obstacle's position using vehicle egomotion parameters. I model the vehicle motion as a forward motion of length s , followed by a change in yaw, γ . The vehicle-relative position of the obstacle is in polar coordinates, where R is the range to the target, and ϕ is the azimuth.

We have the state update equations:

$$\phi_{i+1} = \text{atan}\left(\frac{x}{y-s}\right) - \gamma$$

$$R_{i+1} = \sqrt{(y-s)^2 + x^2}$$

$$x = R_i \sin \phi_i \quad y = R_i \cos \phi_i$$

Equation 6-1. Update equations for the obstacle range R and azimuth ϕ based on vehicle forward motion s and change in yaw γ .

We can also update the uncertainty on these estimates as in the update portion of a Kalman filter:

$$J_1 = \frac{\partial X_{i+1}}{\partial X_i} = \begin{bmatrix} \frac{\partial R_{i+1}}{\partial R_i} & \frac{\partial R_{i+1}}{\partial \varphi_i} \\ \frac{\partial \varphi_{i+1}}{\partial R_i} & \frac{\partial \varphi_{i+1}}{\partial \varphi_i} \end{bmatrix} = \begin{bmatrix} \frac{R_i - s \cos \varphi_i}{\sqrt{k}} & \frac{R_i s \sin \varphi_i}{\sqrt{k}} \\ \frac{-s \sin \varphi_i}{k} & \frac{R_i(R_i - s \cos \varphi_i)}{k} \end{bmatrix}$$

$$J_2 = \frac{\partial X_{i+1}}{\partial Y} = \begin{bmatrix} \frac{\partial R_{i+1}}{\partial s} & \frac{\partial R_{i+1}}{\partial \gamma} \\ \frac{\partial \varphi_{i+1}}{\partial s} & \frac{\partial \varphi_{i+1}}{\partial \gamma} \end{bmatrix} = \begin{bmatrix} \frac{s - R_i \cos \varphi_i}{\sqrt{k}} & 0 \\ \frac{R_i \sin \varphi_i}{k} & -1 \end{bmatrix}$$

where $k = R_i^2 + s^2 - 2sR_i \cos \varphi_i$

Equation 6-2. Jacobians for use in the covariance update on obstacle state parameters R and φ .

The Kalman filter compound equation then tells us:

$$C_{i+1} = J_1 C_i J_1^T + J_2 C_Y J_2^T$$

where $C_Y = \begin{bmatrix} \sigma_s^2 & 0 \\ 0 & \sigma_\gamma^2 \end{bmatrix}$

Equation 6-3. Covariance update equation for obstacle range and azimuth.

After updating, we have an estimate of position and variance of each obstacle's position. Next, the new laser data is classified into non-obstacle and obstacle candidate pixels.

6.3 Classification

The classification module is responsible for finding potential obstacle candidates in each laser scan. Classification is performed using only the current laser data. I first show that intensity information alone can provide better obstacle classification than laser range information in some cases.

At long distances, the returned signal from the road surface is very low power, resulting in unreliable range measurements (see Figure 6-3). In these situations, it makes it impossi-

ble to detect obstacles solely on range. However, the intensity signal provides an easy means of detecting many obstacles (see Figure 6-3).




Figure 6-3. Range and intensity for a single row of laser data. The range (upper trace) has been filtered to remove invalid range points and remove effects of the ambiguity interval and is graphed against the left y-axis. The intensity (lower trace) is unfiltered and graphed against the right y-axis. Even after filtering, range is extremely noisy because of the low signal strength and makes obstacle detection impossible at these distances and this signal level. In contrast, the peak intensity corresponds to the only obstacle in the laser field of view.

Although it is somewhat difficult to see from the graph, the pixels on the obstacle return sufficient signal to provide consistent range estimates. Even with obstacle range information, however, it would be very difficult to segment the obstacle based on the range information since none of the surrounding pixels provide reliable range information with which to compare the obstacle range.

At these distances and signal levels, it appears far easier to look for bright spots in laser scans and segment obstacles based on their intensity relative to the road. The analysis in the previous chapter tells us how much brighter obstacles may be compared to the road. Based on the reflectance model described, we can see that many surfaces, depending on the surface roughness, when seen from a grazing angle may be up to 50% as reflective as when seen from a near-normal direction. A surface should be considered an obstacle if it is twice the intensity or greater than the road surface assuming a similar albedo. Two additional factors make obstacles brighter than the road. First, asphalt has a low albedo compared to

many surfaces. Second, obstacle pixels should be closer to the sensor than road pixels on either side of it. Our reflectance model tells us that closer pixels are brighter -- intensity is proportional to the inverse range squared.

I have tried a number of intensity-based methods to segment obstacles from the background, each of which is described in the following sections. The currently employed segmentation method is described in Section 6.3.3 on page 97.

6.3.1 Road Model-based Thresholding

The first obvious method thresholds on the returned intensity. By driving over an area with no obstacles, it is simple to build a reflectance model for the road including a mean intensity and standard deviation. Thresholding at a reflectance 3 standard deviations above the mean road intensity can detect many obstacles and results in few false positives (approximately 1% of the road pixels are detected as obstacles if the distribution is Gaussian). The false positives can then be eliminated through either minimal obstacle size requirements (wider than a single pixel) or if not confirmed in future scans, can be eliminated as described in the update module.

Unfortunately, this method has a couple of problems. First, the road pixel intensities do not necessarily follow a Gaussian distribution. This can result in segmenting more road pixels as obstacles than expected. Second, the mean reflectance of the road changes fairly rapidly as the sensor pitches relative to the road due to changes in the vehicle suspension causing the laser to hit the road at different ranges. This second problem impacts the detection problem in a number of ways, including guarantees of coverage, etc. Active control of sensor pitch to compensate for vehicle suspension might eliminate some of these problems. However, vertical road curvature will still greatly affect the effective lookahead of the laser sensor (as discussed in Section 3.2.2) and thus the average road reflectance. Figure 6-4 dis-

plays the average and standard deviation of road intensities over time for a single test run (approximately 60 meters of travel at a slow speed in a parking lot).




Figure 6-4. Vertical road curvature or vehicle pitching can cause significant variations in the actual lookahead of the laser sensor. Changes in lookahead affect the mean intensity of the road pixels. The graph illustrates the mean and standard deviation of the road pixels (all pixels not classified as obstacles) at each scanline as the vehicle moves over a parking lot surface. The increase in the standard deviation of intensity at the end is caused by the sensor seeing some white-painted lines marking parking spots.

In the absence of pitch variations or with the addition of pitch control, it might be possible to build and track a slowly-changing road reflectance model over time. In the current system, however, since the average road reflectance can change rapidly due to vehicle pitch variations, I base the segmentation methods on individual scan lines rather than evolving a road reflectance model over time.

6.3.2 Edge-based Segmentation

Edge-based detection methods were also considered. However, large noise levels in the intensity can make edge-based methods susceptible to false positives without significant smoothing (which may obscure narrow obstacles). In addition, statistical methods have the advantage that all points on the obstacle (including non-edge pixels) are automatically selected. Complications in determining the extent of the obstacle can occur in edge-based methods if the edge-detection fails to find one of the obstacle edges.

6.3.3 Histogram-based Segmentation

The current segmentation method is based on histogram splitting (see Figure 6-5). The algorithm finds the end of the histogram peak that contains the median pixel value (I assume that at least 50% of the scanline pixels are road pixels). If the end of the peak is used as the reflectance cutoff, the system would typically detect all objects but also return a significant number of false positives in the tails of the road intensity distribution. To reduce the false positives caused by bright spots on the road we can multiply this cutoff by a “safety factor” (approximately 1.5). Since vertical surfaces typically have reflectances at least twice that of horizontal surfaces, this cutoff still manages to detect almost all objects that pass the other requirements such as minimum width.

Objects with albedos significantly less than the road are typically not detected because they may be statistically similar to the road pixels. If the mean obstacle intensity lies within the distribution of road pixel intensities, the obstacle will not be detected. More laser power might allow us to distinguish between the road and some very dark surfaces by reducing noise and narrowing the distributions, but currently system noise makes this impossible.




Figure 6-5. Intensity histogram for the same row of data shown in Figure 6-3. The segmentation method finds a histogram split point at an intensity around 50. This split-point intensity is multiplied by a safety factor (typically about 1.5) to determine an minimum obstacle intensity.

On rare occasions, depending on the splitting method used, histogram splitting fails to find the maximum road intensity. This may result in false positives, and can occur when the road pixel intensity distribution is multimodal (see Figure 6-6). A multimodal distribution may occur when there is a unusually bright patch on the road. In such cases, the bright feature on the road may be segmented as an obstacle. Unless the feature is very long, however, it will be unverified in future laser scans and will be discarded, as a spurious candidate. Even if the feature is verified in future scans, it may be possible to use the range estimation module (described in Section 6.5 on page 105) to determine whether the area truly belongs to an obstacle or not.




Figure 6-6. An intensity histogram for a row of data with a large bright patch on the road and an obstacle. A significant number of road pixels (to the right of the safety factor cutoff) may be segmented as obstacles because of the bimodal distribution of the road pixels.

Next, obstacle pixels are grouped into individual obstacle candidates in the line. Because the segmentation is statistical rather than feature-based, and because there may be significant noise in the intensity, it is possible that a small percentage of pixels interior to an obstacle may not be classified as obstacle pixels. To reduce the chance that a single obstacle is broken into multiple smaller ones, small non-obstacle areas of 3 pixels or less are reclassified as obstacle pixels. The implicit assumption here is that obstacle distribution is sparse. Then, contiguous obstacle pixels are grouped into single obstacle candidates. The

effect of this gap-filling step is shown in Figure 6-7. These candidates are then merged with existing obstacles on the active list as described in the next section.




Figure 6-7. The bar graph shows a small area of pixels in a single scanline. The height (not drawn to scale) represents pixel intensity. Black bars correspond to pixels which exceed the statistical threshold intensity and are classified as obstacles. Gray bars correspond to pixels which are reclassified as obstacle pixels after gap-filling. Without gap-filling, the object would be broken into 4 small pieces. After gap-filling, only one large obstacle candidate is found (the correct answer).

6.4 Merge

Every candidate found in the classification step is compared to every obstacle in the active list until a suitable match is found. Matching is performed by comparing the predicted azimuth of each obstacle on the active list to the azimuth of the candidate obstacle. If the candidate overlaps with the predicted azimuth of an obstacle, the candidate is matched. The candidate position is merged with the predicted position to provide a refined estimate of obstacle azimuth. Mathematically:

$$\varphi_{i+1} = \varphi_i + \frac{\sigma_{\varphi_i}^2}{\sigma_{\hat{\varphi}}^2 + \sigma_{\varphi_i}^2} (\hat{\varphi} - \varphi_i)$$

$$\sigma_{\varphi_{i+1}}^2 = \frac{\sigma_{\hat{\varphi}}^2 \sigma_{\varphi_i}^2}{\sigma_{\hat{\varphi}}^2 + \sigma_{\varphi_i}^2}$$

Equation 6-4. Merging equations for refining the obstacle azimuth. φ_i is the previous azimuth estimate, and $\hat{\varphi}$ is the azimuth of the matching obstacle candidate from the current laser data.

In addition, all intensity points in the laser data corresponding to the obstacle candidate are added to a list maintained by the active obstacle along with the vehicle position. This information is then used to provide a new range estimate as described in Section 6.5.

Since obstacles are rather sparse, the $O(n^2)$ matching operation takes negligible time. One approach to speed this up, however, would be to sort the active list by azimuth. Matches for candidate obstacles could be found more quickly in this way -- sparse, static obstacles will not likely switch azimuth order.

If no match is found for a candidate, a new obstacle is created and added to the active list. All raw intensities for this object are added to this new obstacle's list. If these obstacles are not verified in subsequent scans, they may be discarded in the update module.

Obstacles are only considered valid if spotted in at least n different scanlines where n is called the "tracking cutoff." Typically we choose $n = 10$. Assuming no significant pitch variations in the vehicle, and that the obstacle is detectable in every line that it is visible, we can relate the height of the object to the number of frames, or scanlines in which it appears (see Equation 3-5 on page 32 and Figure 3-4). This equation implies another trade-off in system design between the minimum obstacle height and the false positive rate. To reduce false positives, we can demand that the obstacle be visible in more lines, but at the cost of not being able to detect very short obstacles.

A sample obstacle-tracking is provided in Figure 6-8. The "image" in the figure is a compilation of individual laser line scans collected as the vehicle moves towards several obstacles. At the start time at the top of the image, the laser only sees the ground of a parking lot. As the vehicle moves, it "pushes" the laser scan along the ground in front of it, generating a "push image."

An overhead view of the environment is illustrated in Figure 6-9. At the beginning (time t_0), the laser sees only ground pixels. At t_1 the laser begins sweeping across a lightpost which results in a stronger laser return (represented as dark pixels in Figure 6-8). As the vehicle moves closer, the vehicle-relative azimuth of the lightpost changes until it finally passes out of the 60° horizontal field-of-view at t_3 . Around time t_2 , the sensor starts

to see several other objects: a second lightpost, a wooden crate, and a parked vehicle. All of these vertical surfaces reflect light more strongly than the road, again represented as dark pixels in the inverted push image. There are also some unusually bright pixels on the ground (dark in the inverted image) surrounding the crate and beyond it (lower in the push image). I believe these bright pixels are caused by parking space markers and perhaps a slight mound in the parking surface.




Figure 6-8. A “push” image. For better printing, the image intensity has been inverted. The image shows a time-series of single line laser scans taken as the vehicle moved in a nearly empty parking lot. At the start (top line), the laser hits only ground pixels. As the vehicle moves forward, the laser is “pushed” along the ground, and sees several obstacles (labelled dark pixels).




Figure 6-9. An overhead schematic of vehicle motion for the laser series in Figure 6-8 (not drawn to scale). The circular arcs represent where the laser scans would hit the ground at each of 4 time instants, beginning at t_0 . The wedge at each position indicates the horizontal field of view. At time t_0 , representing the first scanline in Figure 6-8, 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_3 when it passes out of the field of view. At t_2 , the laser hits the wooden crate near the center of its scan. Shortly after t_2 , the laser hits a parked vehicle and another lamppost.

Figure 6-10 shows the actual segmentation and tracking results from the reflectance data shown in Figure 6-8. After each line is segmented, obstacle candidates are matched to existing obstacles. In the image, individual obstacles are distinguished by different intensities or colors. The obstacles are segmented and tracked properly. Each of the 3 dominant obstacles (2 lampposts and the wooden crate) is detected and tracked as a single obstacle. Segmenting each as a single obstacle is important in obtaining good results from the range estimation module described in the next section. Poor segmentation or a tracking failure

reduces the amount of valid intensity data that can be used coherently to estimate the obstacle range.




Figure 6-10. Segmentation and tracking results from the reflectance data shown in Figure 6-8. Dark pixels are areas segmented as obstacles. Individual obstacles are distinguished by different intensities or colors. Both lampposts and the wooden crate are detected and tracked properly. Specular reflections from the parked vehicle are also segmented and tracked properly.

A parked vehicle at the right periphery is detected and tracked as multiple smaller obstacles. The poor segmentation of the vehicle is caused by its highly specular surface. Some areas on the vehicle reflect most of the laser energy away from the receiver causing a data dropout or a very low intensity pixel. These dropout areas are classified as non-obstacles. Note that the addition of laser range information would not help in these dropout areas since the laser range is invalid here as well. Although these dropouts are clearly undesirable, they do not pose a significant threat. Empirically, there is always a portion of the vehicle that reflects enough laser energy for the vehicle to be detected.

To reduce the clutter caused by false positives, Figure 6-10 only displays obstacles after they have been detected five or more times (which is equivalent to using a tracking cutoff of 5). Nevertheless, there are a few spurious obstacles detected (bottom-center) which appear five times or more. These areas can be attributed to small areas of unusually bright road pixels, some caused by white lines marking parking spots. Since these areas are

detected only briefly, the system may conclude that these “obstacles” are either spurious or very short and thus pose no threat. Using a tracking cutoff of 10 does just that.

Figure 6-11 illustrates the state of the active obstacle list at each frame. The image is similar to the previous figure, but in addition to candidates found in the current laser scan-line, it displays the expected location of obstacles on the “active” list that were not confirmed in the current laser scan. If the vehicle moves several meters without seeing the object, the system removes it from the active list. All of the spurious obstacles detected and illustrated in Figure 6-10 are shown for many frames until they are discarded as flukes. Again, to reduce clutter, Figure 6-11 only shows areas where an obstacle has been detected at least five times. The real obstacles are also moved to the inactive list when the laser can no longer see them (to speed up matching and prevent false matches), but are still considered threats because they were seen in many laser scans.




Figure 6-11. This image illustrates the internal representation of the active obstacle list. In addition to obstacles detected in each line, it shows where the system believes obstacles might be in azimuth at each frame. Even when it can no longer see a potential obstacle, it continues to update the obstacle’s believed position for a short while, based on vehicle egomotion. After moving some distance without seeing it, the object is moved to the inactive list. Depending on where it is located and how many times this inactive “obstacle” was spotted, an avoidance system might ignore it completely, or continue to consider it a threat.

Additional classification and tracking results are provided in Chapter 7. The system was able to detect many obstacles at distances up to 60 meters while eliminating nearly all false positives.

6.5 Range Estimation

While it is not necessary to have centimeter-accurate range estimates to highway obstacles, it would be useful to have range estimates within a few meters. A measure of uncertainty would also be useful since the uncertainty in the range estimate imposes uncertainty on the predicted azimuth of the obstacle in the next frame.

Because of vehicle pitch variations and vertical curvature of the road, range estimates based on sensor inclination are probably only accurate to within 20 meters. As we saw in the previous chapter, laser reflectance is dependent upon range. We can use this dependence to our advantage to estimate an object's range based upon the returned intensity. Since intensity is also dependent upon other unknown factors, a single intensity reading will not provide us with range information. However, by using egomotion estimates of the vehicle and by tracking the obstacle over multiple frames, we can estimate the range.

The reflectance model in the previous chapter shows that intensity depends on many parameters including albedo, surface shape, surface orientation, surface roughness, and a specular ratio. Given the small amount of data present for each obstacle, it is nearly impossible to solve for all of these parameters. It is also unnecessary. Once pixels are classified as an obstacle, the only parameter we care about is the range. If we knew the other parameters, they would be useful in range estimation, but they are unimportant for planning or control for obstacle avoidance in and of themselves. All the parameters except range and orientation are constant for a given obstacle. Since our required field-of-view is fairly small (only $\pm 20^\circ$), orientation changes (up to the size of our field-of-view) as we approach a static obstacle will have relatively small effects on the reflectance except for highly specular objects. Range changes as the vehicle approaches the obstacle will have a dominant

effect on the reflectance, particularly for those obstacles which are nearly straight ahead. These constraints allow us to estimate range based on a much simpler model:

$$IR^2 = \lambda$$

Equation 6-5. Ignoring changes in orientation (which will be small), the product of the intensity I and the range squared, R , is equal to a constant λ , called the normalized reflectance.

Since the vehicle did not have a compass or other accurate means of estimating yaw, the vehicle was driven along a straight path as much as possible. A straight-line path was then assumed for each of the range estimation methods. The methods could be altered to account for non-linear motion by maintaining more state information.

My first attempt to estimate range from intensity used an extended Kalman filter as described in Appendix B. The normalized intensity, obstacle range, and obstacle azimuth were the state parameters for the filter. At each step, the obstacle range and azimuth were updated based on the vehicle motion to predict the new location of the obstacle in the image. The normalized reflectance was unchanged. Next, a new laser measurement was taken which provided new intensity information and a new estimate of the obstacle azimuth. The range and normalized intensity state parameters were updated using a form of the relationship given in Equation 6-5.

Unfortunately, the Kalman filter failed to estimate the range correctly because the error between estimates and predictions in the merging step were correlated over time. Even with very large initial uncertainties, the filter became relatively certain of the wrong answer because at each step it arrived at nearly the same wrong answer. These correlated errors indicated that an iterative solution was unlikely to work.

To enable a memory-based range estimation algorithm, I modified the obstacle tracker to remember all previous intensity measurements with vehicle position data (assuming linear travel) for each active obstacle. At each time step, the estimator uses all previous data

to provide a current range and a range-normalized intensity estimate. Figure 6-12 illustrates the relationship between position changes and the obstacle range.

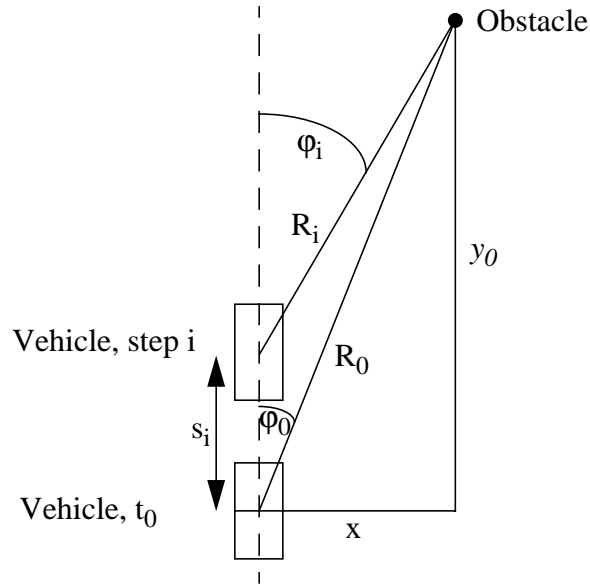


Figure 6-12. The above figure illustrates the relationship between the distance to an obstacle at time step i , R_i , and the initial range R_0 . This assumes that the vehicle travels in a straight line such that s_i is the travel distance at time t_i since the time t_0 when the obstacle is first discovered. ϕ_0 and ϕ_i are the azimuthal positions of the obstacle at time t_0 and t_i , respectively.

Mathematically, we relate the current range R_i to the initial range R_0 through the travel distance of the vehicle and the azimuth of the obstacle at each time step:

$$R_0 \cos \phi_0 = y_0$$

$$R_i \cos \phi_i = y_0 - s_i$$

Substituting the above expressions for the range in Equation 6-5 on page 106, we can relate the current intensity to the normalized reflectance, the travel distance, the initial longitudinal distance, and the current azimuth.

$$I_i \left(\frac{y_0 - s_i}{\cos \phi_i} \right)^2 = \lambda$$

Equation 6-6. The normalized intensity λ of a target is equal to the product of the noise-compensated intensity and the square of the range to the target. y_0 is the initial longitudinal range, s_i is the longitudinal distance travelled, and ϕ_i is the target azimuth in the laser image.

To estimate the current range to an obstacle, we need to estimate the two unknowns from Equation 6-6: y_0 (the initial longitudinal range) and λ (the normalized intensity), based on all reflectance and position data for an individual obstacle.

I tested 5 statistical methods for estimating y_0 and λ , described in Section 6.5.1 through Section 6.5.5. All of the range estimates are based on some form of minimizing the residuals between the predicted intensity as dictated by the reflectance model in Equation 6-6 and the measured intensity. Section 6.5.7 discusses generating uncertainty estimates on the range computed by the 5 statistical methods. An additional modification is made to account for intensity noise since it has a nonzero mean equal to the mean photon noise as described in Section 5.2.1. Mathematically, the residual can be stated:

$$r_i = I_i - \mu_{PN} - \frac{\lambda \cos^2 \phi_i}{(y_0 - s_i)^2}$$

Equation 6-7. The residual is the pixel intensity minus the predicted intensity given position data ϕ_i and s_i and the model parameters λ (the normalized intensity) and y_0 (the initial longitudinal range). I compensate for the nonzero mean noise in the intensity by subtracting the mean ambient photon noise μ_{PN} from the measured intensity. The ambient photon noise can be measured by a scan prior to the start of the experiment.

Since the residual is a nonlinear function of the data, linear solution methods can not be used. Fortunately, there is a closed-form solution for a least squares estimate of λ in terms of y_0 (see Equation 6-8). Unfortunately, the expression for λ is not easily separable in terms of y_0 , so the summations for computing λ must be recomputed for each proposed y_0 . Then a search is performed for the optimal y_0 (according to one of the 5 statistical methods described below) over the range 20 to 80 meters at a one meter resolution.

$$\frac{\partial}{\partial \lambda} \sum_i^N r_i^2 = 0 \Rightarrow \lambda = \frac{\sum \frac{(I - \mu_{PN}) \cos^2 \phi_i}{(y_0 - s_i)^2}}{\sum \frac{\cos^4 \phi_i}{(y_0 - s_i)^4}}$$

Equation 6-8. Closed form least-squares solution for the normalized reflectance λ as a function of y_0 , the initial longitudinal range.

If we put the residual in terms of λ rather than I , and use Equation 6-9, we arrive at alternative formula for λ that is separable in terms of y_0 (see Equation 6-10). Although this new formula for λ does not provide the same answer for λ for all y_0 , it does arrive at the correct

solution for the best y_0 . This is enough to achieve nearly the same range estimates, and this formula for λ can be computed in constant time if a few running sums are kept as new data is added.

$$r_i = \frac{(I_i - \mu_{PN})(y_0 - s_i)^2}{\cos^2 \phi_i} - \lambda$$

Equation 6-9. An alternative formulation for the residual. Used only for the computation of λ , not for y_0 .

$$\frac{\partial}{\partial \lambda} \sum_i^N r_i^2 = 0 \Rightarrow \lambda = y_0^2 \sum k_i - 2y_0 \sum k_i s_i + \sum k_i s_i^2 \quad \text{w where } k_i = \frac{I_i - \mu_{PN}}{\cos^2 \phi_i}$$

Equation 6-10. An alternative formulation for λ that is computable in constant time for different y_0 if 3 running sums are kept for each of the summation terms.

Although it is acceptable to use the formulation for the residual in Equation 6-9 to arrive at Equation 6-10 for the calculation of λ , we must still use the formula given for the residual in Equation 6-7 for calculating y_0 , since we want to minimize the error between our measurement (the intensity) and its value as predicted by the parameters λ and y_0 . If we had a good measure for the standard deviation of the error in the intensity, σ_I , we should properly divide the residual by σ_I .

6.5.1 Least Squares

Ordinary least squares regression techniques estimate desired parameters by minimizing the sum of the squared residuals. Mathematically:

$$\text{minimize } \sum_{i=1}^n r_i^2$$

Equation 6-11. Least squares estimation chooses the parameter set which minimizes the sum of squared residuals, r_i .

Although least squares methods are among the most commonly used methods in parameter estimation, they perform very poorly with respect to outliers. A single outlier can have an arbitrarily large effect on a least squares estimate. Thus, I designed and tested 4 more methods based on robust statistics.

6.5.2 Least Median of Squares

The first robust method tested was a least median of squares (LMS). There are many robust statistics techniques that can produce estimates that are unaffected by some outliers. One way to characterize a statistical method is by its breakdown point ϵ^* , which is the smallest percentage of outliers that can result in arbitrarily large errors in the resulting estimate. Whereas least squares has a breakdown point of 0%, the least median of squares technique developed by Rousseeuw has a breakdown point as high as 50%, the maximum possible[44]. Least median of squares can be stated:

$$\text{minimize } (\text{median}(r_i^2))$$

Equation 6-12. Least median of squares estimation chooses the parameter set which minimizes the median of the squared residuals, r_i .

Typically, a least median of squares approach searches over all parameter sets that may be defined by minimal subsets of the data points. For example, to find the model of a line $y = mx + b$ which minimizes the median of squared residuals, a search is performed over the models (m, b) defined by every unique pair of points (x, y) in the data set. If there are many points, this can result in a large number of models to test.

Instead, I search over a limited range around the least-squares solution to λ and over the limited domain for y_0 mentioned previously. This reduces the number of tested models significantly.

6.5.3 Least Trimmed Squares

Another robust statistic method developed by Rousseeuw is “least trimmed squares” (LTS). For every parameter combination, the residuals are calculated for every point and then sorted. The h smallest squared residuals are summed, and the parameter combination which minimizes this sum is chosen. Mathematically:

$$\text{minimize } \sum_{i=1}^h r_i^2$$

Equation 6-13. Least trimmed squares estimation chooses the parameter set which minimizes the sum of the h smallest squared residuals, r_i .

A search over the same parameter space used in the LMS estimate was used for the LTS estimate. Note that calculating the LTS is very expensive since each sort operation requires $O(n \log n)$ operations.

In my experiments, I found that the LTS method (using $h/n = 0.5$), despite its greater computational expense, performed similarly to LMS. As a result, I stopped using LTS for range estimation.

6.5.4 Least Squares of Maxima

Another method for range estimation tracks the maximum obstacle intensity at each frame and then uses only these maximum intensity measurements to generate a least-squares fit to the model.

There are a couple advantages to using this method. First, the area of maximum intensity will be the area that is most reliably and consistently segmented as an obstacle. As such the maximum is the measure least likely to be affected by segmentation errors. Except for highly specular surfaces where small changes in position will drastically affect the returned intensity, the maximum intensity at each time instant should correspond to the same point on the surface, i.e. the point with maximum albedo.

There are also two problems with this method, however. First, the maximum may be quite susceptible to noise since it may find measurements out on the tail of the noise distribution. Second, because the laser spot size changes with distance, this method can sometimes underestimate the distance to an object. At long distances, the instantaneous laser spot may be larger than brightest surface patch and will average darker areas in the measurement. At near distances, the laser spot may be small enough that it falls entirely within the bright area on the surface. This will cause the intensity to climb more rapidly than predicted by the inverse-square law as the vehicle draws closer and will cause the method to underestimate the range.

6.5.5 Least Squares of Medians

This method is similar to the previous method. Instead of taking the maximum intensity at each time instant, it takes the median, and uses only these median intensities to calculate the least squares parameter set.

The median has the advantage of being a fairly robust measurement. However, low albedo areas on the object surface are less likely to be segmented as obstacle at long distances, and hence the median may be corrupted somewhat by poor segmentation. Errors of this type will typically generate overestimates of the range since darker areas may be thrown out at longer distances resulting in artificially high medians, but will be averaged in at near distances. The intensity will rise more slowly than expected, causing the overestimate.

6.5.6 Range Estimation Results

While range cannot be estimated reliably for objects that are tracked only briefly (such as the parked vehicle), reasonable range estimates can be made for the two lightposts and the wooden crate from the tracking results shown in Figure 6-8 through Figure 6-12.

Range estimation results for the first lightpost are given in Figure 6-13. When the laser first intersects the lightpost it is aimed approximately 35 meters ahead of the vehicle. Intensity-based range estimates are essentially random at first. As the vehicle travels towards the object and gathers more intensity information, the range estimates begin to converge to the correct range. All of the range estimates provide reasonably close answers to the actual range (as measured by the laser range channel) after the vehicle has moved approximately 10 meters towards the object.

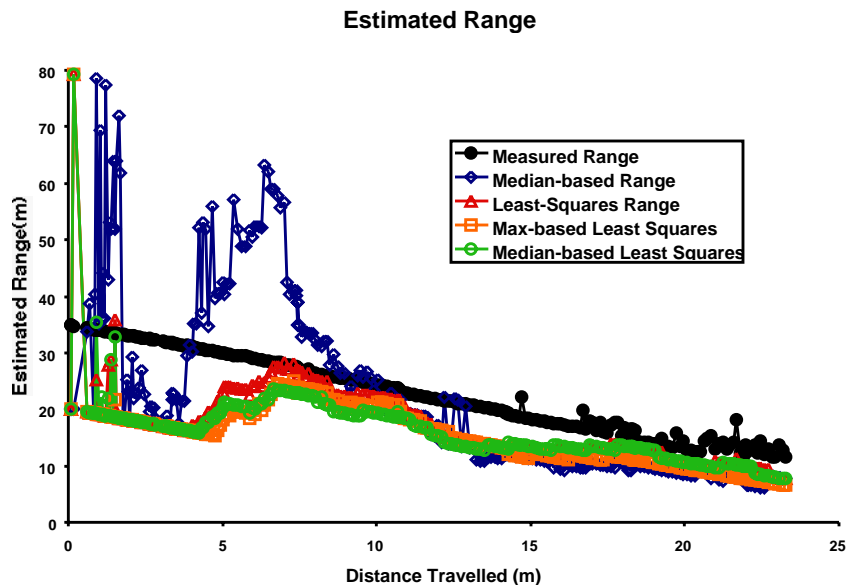


Figure 6-13. Range estimation results for the first lightpost from the earlier tracking examples. Range estimates are essentially random at first, but as the vehicle travels and gathers more intensity and position data, it begins to provide good range estimates. After 10 meters of travel, all of the range estimates are in close agreement to the actual range (as measured by the laser range channel). All of them underestimate the actual range, but are mostly within 5 meters. The estimates drift away from the actual range for several meters before returning to the actual answer.

Figure 6-14 provides range estimation results for the wooden crate.

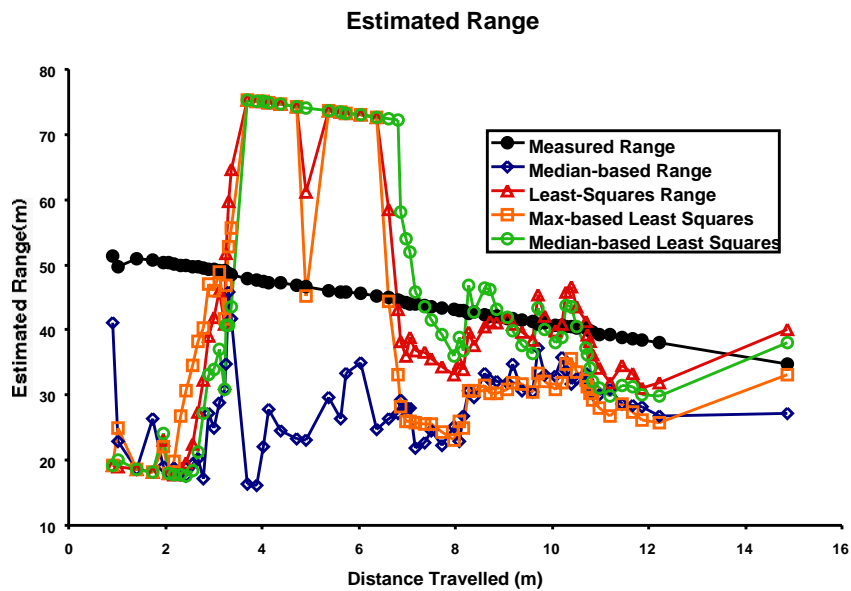


Figure 6-14. Range estimation results for the crate. The least squares and the least squares of medians methods provide reasonable range estimates after 10 meters of travel. After approximately 12 meters of

travel, the crate disappears under the laser field-of-view. The last detection at 15 meters is caused by a false positive.

Figure 6-15 provides range estimation results for the second lightpost.




Figure 6-15. Range estimates for the second lightpost. Again, after 10 meters of travel all range estimates begin to lock in on the correct range.

The range estimation results for the lightposts bear striking similarities after about 10 meters of travel. The similarities suggest that the range errors are not caused by intensity noise; errors caused by noise should be random. Instead, many of the errors are likely caused by position errors, intensity drift, or other real phenomena. The effectiveness of the range estimation for the crate is more difficult to judge since it passes below the laser field-of-view once the estimates converge near the correct range. This is unavoidable with short obstacles and a single line VFOV. For this reason, additional experiments were conducted with an extended vertical field of view, and are discussed further in Chapter 7.

All graphs show that the 4 methods begin to provide accurate range estimates after about 10 meters of travel. In each case, range estimates drift away from the actual range for several meters before drawing close again. All graphs also show the methods mostly underestimating the actual range after convergence. For the lightposts, the drift of the range estimates away from the actual range occurs over approximately 7 meters of travel in each

graph. These drift errors may be caused by optical crosstalk which is discussed further in Section 6.6.

The underestimation in the range could be due to noise, unmodeled reflectance phenomena, or a bias in the vehicle position estimate. If the vehicle position sensor underestimated the actual distance travelled, it would cause underestimation of the obstacle range. Another possible cause for the bias is a drift in intensity. If intensity measurements drifted upwards, the range would be underestimated. However, as seen in Section 5.4, intensity drift is generally in the opposite direction (negative), and it is slow enough that is unlikely to cause problems in this particular experiment which took approximately a minute to complete. The additional range estimation experiments in Chapter 7 seem to indicate that the estimates eventually converge to the correct range for diffuse surfaces. Specular surfaces are more problematic, however.

6.5.7 Estimating Uncertainty in Range

Noise in intensity measurements will cause uncertainty in the range estimates. Unfortunately, it is difficult to provide a precise estimate the uncertainty on the range estimates provided by any of the statistical methods described above. The search-based range estimation technique does not provide a direct means of estimating the uncertainty. The sum of the squared residuals in the least-squares technique for each proposed y_0 , however, does provide us some information. When the software first begins to track an obstacle, all of the proposed y_0 are equally valid and the sum (or median in the LMS method) of the squared residuals for each y_0 are similar. Later, as more intensity measurements at multiple ranges become available, the sum (or median) of the squared residuals is much larger for y_0 far from the correct y_0 (see Figure 6-16 and Figure 6-17).




Figure 6-16. The evolution of the mean squared residual over time for multiple proposed y_0 for the first lightpost in Figure 6-8. When the system first begins tracking an obstacle, there is little information, and all potential y_0 yield similar mean squared residuals. At the end, the system converges to a y_0 which yields the minimum mean squared residual, e.g. a y_0 of approximately 40 meters (the correct answer).

Uncertainty can be estimated by finding the minimum and maximum y_0 which have an error “close” to the error on the best y_0 . Determining how “close” an error should be to the best error to consider its corresponding y_0 as possibly valid is difficult, however. “Close” is also likely to be different for each of the error metrics.




Figure 6-17. The evolution of the median squared residual over time for multiple proposed y_0 . When the system first begins tracking an obstacle, all potential y_0 yield nearly equivalent median squared residuals. Similar to the least squares method, the least median of squares converges over time to the correct y_0 .

6.6 Problems and Complications

The reflectance based obstacle detection method described in this chapter has problems both inherent and advenient to the method, some of which we address in this section.

A problem inherent to the intensity-based method is that a given intensity measurement has multiple possible causes or interpretations. A pixel of medium brightness may be caused by a white object at a long distance or a dark object at a close distance or an infinite number of other possibilities. This many-to-one mapping from the 3D world to the 2D image is a problem inherent to all vision methods that try to reconstruct 3D world information. Assuming a relatively flat world populated by some obstacles, as in a highway environment, removes some possibilities, but does not eliminate all ambiguity. In particular, a dark obstacle may appear like the road, and a bright road patch may look like an obstacle.

6.6.1 Dark Obstacles

The current system is unable to detect objects with albedos much lower than the road. Road pixel intensities often have near zero intensity at long ranges (see Figure 6-5 on page 97 for an example histogram), which makes them inseparable from darker objects with the current laser power. At short ranges or with increased laser power, black obstacles might be separated from road pixels and detected as obstacles. For example, close range laser scans can produce histograms similar to that in Figure 6-18.




Figure 6-18. Intensity histogram for a single laser scan aimed at the ground approximately 20 meters away. The scan hits several obstacles, some darker than the ground and some brighter than the ground. In cases like this where the road pixels do not reach zero intensity, it may be possible to segment obstacles which are darker than the road. The type of obstacles that can be detected depend on the width of the road pixel distribution.

The main lobe in the histogram in Figure 6-18 corresponds to the road pixel intensities. The types of obstacles that can be segmented in such a scan depends on the width of the road pixel distribution. The distribution width is dependent on both the intensity noise and the actual intensity variations due to differences in local road albedo and roughness. A good model of the road pixel intensity distribution combined with the laser reflectance model developed in Chapter 5 would allow us to describe those obstacles which would “spoo” the road and would be inseparable from road pixels by statistical methods. Although such

an analysis would allow us to predict which obstacles could be confused for road pixels, we can never rid ourselves of the potential confusion for this set of objects.

Dark objects undetectable using the intensity-based methods are usually also undetectable with range measurements, since dark surfaces at long distances typically provide unreliable range measurements. Since range measurements on the road surface are also unreliable, these objects are not easily detectable with range.

6.6.2 Ground Markings as False Positives

Assuming a rectangular beam cross-section, the laser illuminates a trapezoidal area on the ground. An elliptical cross section illuminates a skewed elliptical patch on the ground (see Figure 6-19). The major axis (in this particular geometrical case) of the “elliptical” patch is equivalent to the altitude of the trapezoid, and the minor axis is equivalent to the average width of the trapezoid. Given the beam divergence and the height and lookahead distance of the laser, we can calculate the laser footprint on the ground. The laser beam is elliptical and is supposed to have a Gaussian intensity distribution over both axes of the laser ellipsis. Divergence is defined by the angle where the Gaussian distribution reaches $1/e^2$ (13.5% of the peak value) attenuation. According to laser collimator data sheet, the beam divergence should be 0.1 mrad vertical and 0.3 mrad horizontal. In fact, since the distribution is only Gaussian at short distances, the divergence was also measured by examining the visible spot size on the wall. The eye measures a larger spot by seeing more than the $1/e^2$ attenuation boundary and measures a beam divergence of 0.3 mrad vertical and 1.6 mrad horizontal[36].




Figure 6-19. If we model the beam cross-section as a rectangle, the laser illuminates a trapezoidal area on the road ahead (gray area). The altitude of the trapezoid, or length of the spot, x is determined by the beam inclination angle (or lookahead distance L and height h) and the vertical beam divergence δ_v . The average width of the trapezoid is determined by the lookahead distance and the horizontal beam divergence δ_h . An elliptical cross-section illuminates a distorted elliptical patch on the ground.

Using the eye-based divergence measurements and $h = 1$ m, and $L = 60$ m and Equation 6-14, we find the laser footprint length is 1.08 meters and the spot width is 9.6 cm.

$$x = 2 \left(\frac{h}{\tan(\alpha - \delta_v/2)} - L \right) \quad \alpha = \text{atan}(h/L)$$

$$y = 2L \tan(\delta_h/2)$$

Equation 6-14. The length of the laser spot (along the road) x is determined by the height h of the sensor, the lookahead distance L , and the vertical beam divergence δ_v . The width of the laser spot (across the road) y is determined by the lookahead distance and the horizontal beam divergence δ_h .

This implies that small discolorations or markings on the ground will simply be averaged into the background because of the large spot. Only obstacles or very long longitudinal ground features (such as lane markings) will be distinguishable in the laser image. If a lane marking is initially segmented as an obstacle (because of its relative brightness), we expect that it can be discarded as a threat through 2 types of information. First, a road follower should mark the same region as belonging to a lane marker. Second, unlike an obstacle which is getting closer, the intensity of the lane marker will remain relatively constant over time. We also note that even if a stray ground marking is selected as an obstacle, it will quickly pass out of the laser field of view, and will thus not be confirmed as an obstacle.

6.6.3 Optical Crosstalk

Although the last two problems discussed can never be fully solved, I now discuss one that could be solved through a redesign of the sensor system. Significant optical crosstalk (as described in Section 4.2.2) can make an obstacle appear temporarily dark, resulting in a missed detection or poor range estimate. Figure 6-20 shows a push image with significant and obvious optical crosstalk caused when the plexiglass environmental cover is used. As the vehicle moves closer to the objects, the intensity should increase monotonically in accordance with the intensity versus range relationship described by Equation 6-7 on page 99. Rather than monotonically increasing, however, the intensity recorded on the deer hide and cinder block obstacles oscillates over time. The distance travelled between the local minima or maxima intensities is approximately 6.5 meters, which is the ambiguity interval of the high frequency channel, and is indicative of optical crosstalk as explained in Section 4.2.2.




Figure 6-20. The plexiglass cover increases the effect of range/reflectance crosstalk. The effect of crosstalk can be seen in both two objects in this push image (a compilation of single-line laser scans over time). Intensity should increase monotonically as the vehicle approaches the obstacles. Instead, the intensity of both the deer hide and the cinder block obstacles oscillates significantly with a period of 6.5 meters. This is the same as the ambiguity interval of the high frequency modulation. Laser intensities have been inverted for better printing.

While crosstalk typically does not have visible effects on the measured obstacle intensities when the sensor is uncovered, it does have an effect on our intensity-based range estimation. The effect of optical crosstalk can also be seen in many of the performance graphs

for the range estimation algorithms in Chapter 7. Many of the graphs show some oscillation in the range estimate. The oscillation has a period of approximately 6.5 meters of travel which is again indicative of optical crosstalk.

Fortunately, crosstalk can be eliminated. Crosstalk is caused by the AM modulation of the laser signal. Since AM modulation is required only for the operation of laser range channel, and not for the reflectance channel, it could be eliminated. An unmodulated intensity sensor would be free of this crosstalk problem, and might provide better obstacle detection and intensity-based range estimates in a shorter amount of tracking time and distance.

6.7 Chapter Summary

The obstacle detection system uses a state-of-the-art laser scanner configured in a single line scan mode and a vehicle position sensor. The tracking software consists of 4 components. An update module updates vehicle-relative obstacle positions based on vehicle motion. The segmentation module uses histogram analysis to find obstacle candidates. These candidates are matched and merged into previous obstacle candidates if possible. Finally, the intensity and position estimates are used to estimate obstacle range via memory-based statistical methods.

Results appear promising, and the next chapter shows system results in more detail. Nevertheless, there are a number of problems with the system. Vehicle pitch variations can cause the system to miss an object entirely. Obstacles with very low albedo are not detected. Severe optical crosstalk can also cause missed detections. Finally, specularities and optical crosstalk can cause errors in range estimation.

Fortunately, many of these problems can be solved through better engineering of the laser system and the software. Other problems may be “solved” by using an additional obstacle detection method. For example, dark surfaces could be detected by a stereo system[50]. The next chapter provides more experimental results of the detection, tracking, and ranging system.

CHAPTER 7 **Experimental Results**

This chapter provides additional obstacle detection, tracking, and range estimation results to those presented in Chapter 6.

7.1 Detection and Tracking Results

I performed tracking experiments both indoors and outdoors. Each experiment consisted of 500 to 1500 single-line scans taken over approximately 50 meters of vehicle travel. All experiments listed in this chapter were performed without the presence of the plexiglass environmental cover discussed in Section 4.1.2 on page 54 unless otherwise noted. In addition, all experiments used the classification method described in Section 6.3.3 on page 97, using a “safety factor” of 1.5 except where otherwise noted.

Even with a safety factor of 1.5, spurious detections in individual scans occur with reasonable frequency. Thus, real obstacles are distinguished from false positives by the number of scans in which they can consistently be detected and tracked. We can choose a “tracking cutoff” n , where an obstacle candidate is considered an actual obstacle if it is detected and tracked in at least n laser scans, and is ignored otherwise. The tracking cutoff

is a trade-off between rejection of spurious candidates and real obstacle detection, and the cutoff should be chosen in relation to Equation 3-5 on page 32. An obstacle candidate is considered a “false positive” if it is detected and tracked at least n times (it exceeds the tracking cutoff) and is not attributable to a real obstacle. A missed detection or “false negative” occurs if a real obstacle fails to be detected at least n times (does not meet the tracking cutoff).

7.1.1 Indoor Experiments

I performed a series of experimental runs in the highbay of the Robotics Engineering Consortium. In each case, the laser lookahead distance was between 30 and 50 meters. At the time of these experiments, I did not have the position sensor on the vehicle. Since vehicle position changes are normally used for predicting changes in obstacle location from one scanline to the next, the tracking algorithms were modified to search near the last known obstacle position for a match.

The highbay presented a very cluttered environment. There were dozens of obstacles present, mostly along the walls of the highbay. In each experiment, I placed one obstacle in the center of the vehicle path. Detecting the placed obstacle was the primary objective. In 28 experiments, the placed object was detected every time. The obstacles tested included: a cinder block, a wooden crate, a cement surface, a deer hide, a traffic cone, and a white painted surface. At least two experiments were performed with every object, and every object was able to be detected at a minimum of 35 meters.

I also examined the reflectance data visually to look for false positives and to see if the system failed to detect any of the ambient obstacles. Using a tracking cutoff of $n=10$, I saw no obvious false positives in any of the experiments. Except in the case of very dark obstacles which had albedos significantly less than the ground, I spotted only two misses (false negatives). Both of these false negatives occurred when more than 50% of the horizontal field of view was occupied by obstacles (which violates a critical assumption in the classification module). Since I do not have an independent model of the environment with which to compare the segmentation and tracking results, I cannot be certain of this result.

7.1.2 Outdoor Experiments

Outdoor experiments were all performed in Morewood parking lot at Carnegie Mellon. In each of the outdoor experiments, I used a 60 degree horizontal field of view (HFOV). This is a larger HFOV than necessary for highway operation, but was used to allow tracking of more obstacles than might have been otherwise possible.

7.1.2.1 Experiments without Position Information

Table 7-1 shows detection and tracking results for a series of runs performed in the Morewood parking lot at CMU without vehicle position information. The first column gives an experiment identification number. The second column provides the actual number and type of obstacles present within the HFOV at some time during the experiment. *Total candidates detected* indicates the total number of distinct obstacle candidates that were detected in one or more scans over the course of the experiment. Column four contains 3 numbers for each experiment, indicating the number of candidates that were seen and tracked over at least 5, 10, and 20 times. These numbers are often far less than the total number of candidates because spurious detections or minor tracking errors inflate the total candidate count.

The numbers in column four are also typically greater than the number of real obstacles seen in the actual scans. This does not necessarily indicate that there are any false positives. Individual obstacles may be spatially segmented into multiple objects or temporally tracked as multiple objects for a variety of reasons. Specular surfaces in particular present significant challenges to the tracking algorithm. First, specular or highly heterogeneous obstacles are often split spatially by segmentation since some areas of the obstacle may reflect light away from the laser detector. Second, temporal tracking depends on good azimuth predictions, and azimuth predictions depend on range estimates which tend to be poor for specular objects. Third, specular reflections can appear or disappear quickly based on small changes in vehicle position. A Bayesian tracking algorithm which could support multiple hypotheses might handle these challenges better.

Columns five and six indicate the actual number of false positives or misses (false negatives) if tracking cutoffs of 5, 10, and 20 are used. In almost all experiments, there are no misses and no false positives using any of these cutoffs. Short obstacles are more likely to be missed, however, if a large tracking cutoff is used.

Since I do not have an independent model of the environment with which to compare my detection results, the numbers in the false positives and misses columns are based on my visual perusal of the reflectance data. I do not require that the entire object surface be classified properly to count as a detection. I only require that some area on the object is detected consistently enough to pass the 5, 10, or 20 line tracking cutoff as this would provide an automated vehicle enough of a warning to avoid the obstacle (assuming some reasonable safety buffer is used). As such, the results are somewhat subjective.

The final column lists the lookahead distance of the laser. Most obstacles are initially detected at approximately this range. A tracking cutoff of 5 is sufficient to eliminate all false positives and allow all obstacles to be detected and tracked correctly.

Table 7-1. Outdoor Segmentation and Tracking without Position Data

Experiment	Number and Type of Actual Obstacles	Total candidates detected	Candidates seen 5/10/20 times	False Positives 5/10/20	Misses 5/10/20	Range (meters)
1	3 total: 1 cinder block, 2 lampposts	29	7/7/6	0/0/0	0/0/0	35
2	3 total: 1 cinder block, 2 lampposts	37	7/6/6	0/0/0	0/0/0	32
3	3 total: 1 cinder block, 2 lampposts	34	8/6/6	0/0/0	0/0/0	33
4	3 total: 1 traffic cone, 2 lampposts	77	8/8/6	0/0/0	0/0/0	32
5	6 total: 1 crate, 2 lampposts, 3 cars	133	48/40/31	0/0/0	0/0/0	60
6	3 total: 1 crate, 2 lampposts	64	7/7/7	0/0/0	0/0/0	33

Obstacle segmentation was relatively trivial in each of the experiments listed in Table 7-1. At a 35 meter lookahead, the laser reflectance signal is strong enough that the

intensity histogram separates obstacles from almost all road pixels. For all of the experiments with a lookahead of 35 meters or less, there are relatively few spurious obstacle candidates, and they are easily filtered out through temporal tracking. The only tracking errors are caused by a processing delay in the data collection process every fiftieth scan (see Figure 7-1). Without odometry information to help predict the location of the obstacle in the new scan, the pause in the data collection can cause the obstacle in the next line to be outside of the candidate tracking window since the vehicle continues to move. The system responds by occasionally creating a new obstacle candidate after a pause instead of matching the new data to the old obstacle. Eliminating the processing delay or enlarging the tracking window would fix this problem, but the addition of odometry information fixes this problem and allows for range estimation as described by the procedure in Chapter 6.

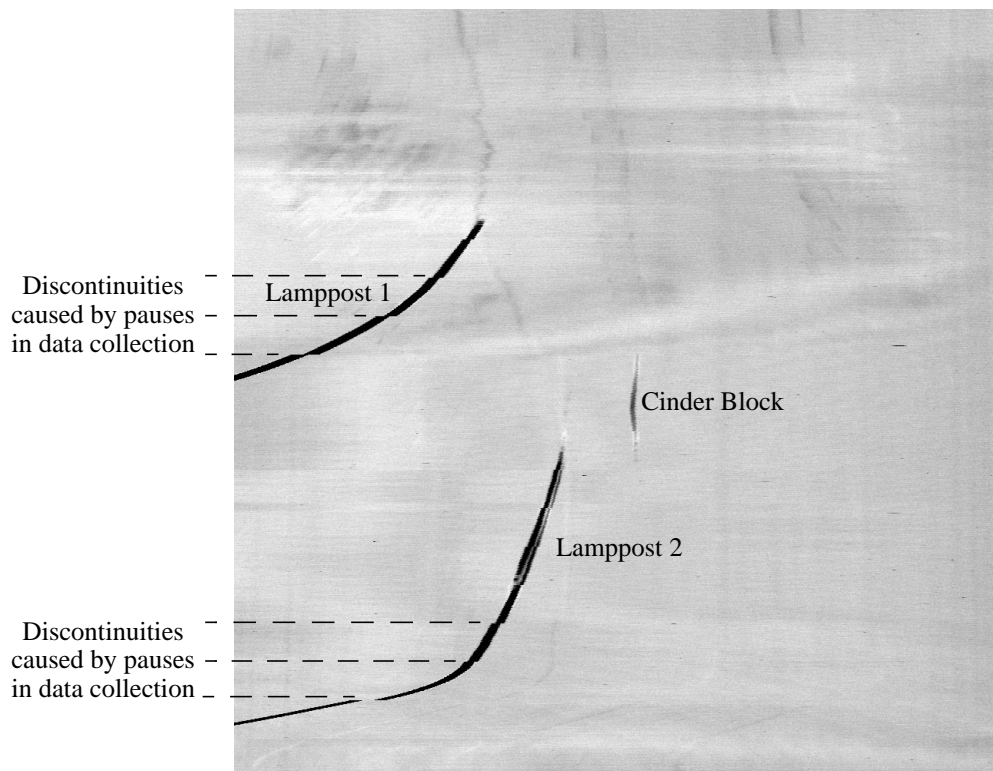


Figure 7-1. A push image for experiment 1 from Table 7-1. Pixel intensities have been inverted for better printing. Segmentation and tracking of the 3 obstacles is accomplished correctly for the most part, except at the discontinuities caused by processing delays in the data collection.

7.1.2.2 Outdoor Position-Tagged Experiments

Table 7-2 shows detection and tracking results for a series of runs performed in the Morewood parking lot at CMU. Every fifth line of laser data was tagged with position information provided by a Datron DLS-1 optical correlation odometry sensor. Position estimates for the remaining scans were calculated by linear interpolation between position readings.

Table 7-2. Tracking with Position Data, without Environmental Cover

Experiment	Number and Type of Actual Obstacles	Total candidates detected	Candidates seen 5/10/20 times	False Positives 5/10/20	Misses 5/10/20	Range (meters)
7	4 total: 1 crate, 2 lamp-posts, 1 car	92	11/9/6	0/0/0	0/0/0	50
8	4 total: 1 crate, 2 lamp-posts, 1 car	85	21/10/8	0/0/0	0/0/0	50
9	5 total: 1 cement surface, 2 lampposts, 2 cars	161	67/38/20	0/0/0	0/0/0	50
10	6 total: 1 cement surface, 2 lampposts, 3 cars	173	72/46/32	0/0/0	0/0/0	60
11	7 total: 1 cement surface, 2 lampposts, 4 cars	146	56/40/25	0/0/0	0/0/0	52
12	5 total: 1 cinder block, 2 lampposts, 2 cars	227	66/32/12	many/7/0	0/0/1	46
13	6 total: 1 cinder block, 2 lampposts, 3 cars	154	67/40/27	0/0/0	1/1/1	58
14	5 total: 1 crate, 2 lamp-posts, 2 cars	172	51/27/17	3/0/0	0/0/0	48
15	5 total: 1 crate, 2 lamp-posts, 2 cars	137	55/37/26	0/0/0	0/0/0	65

The table shows that at lookahead distances of 45 to 65 meters, obstacle detection is still successful in most cases. Most objects were detected and most experiments had no false positives. However, there were a few errors in these experiments which I now discuss in more detail.

Experiment 12 Details. A patch of unusually bright ground pixels causes multiple small false positives in a single area. I believe the bright pixels were caused by multiple parking

space markers, although this is difficult to confirm with only the 1-D laser data. Because of the large laser spot size on the ground, it is difficult to see such markers clearly in the bright pixels. The false positives are eliminated if the tracking cutoff is set at 20. However, this also eliminates an actual obstacle. A cinder block was tracked for over ten lines, but less than 20 lines. The significant height of the cinder block would normally have made it appear in more scan lines, but pitching motion of the vehicle caused the cinder block to pass out of the vertical field of view (VFOV) prematurely. Similar bright patches appear in many of the other experiments. If these patches are caused by white-painted parking space markers as suspected, then this problem is peculiar to this environment and should be absent in a real highway environment. This might allow the use of a smaller safety factor which would improve detection of smaller objects at long distances such as the cinder block.

Experiment 13 Details. A cinder block obstacle at 58 meters is not detected using the default safety factor of 1.5. However, it is detectable if a safety factor of 1.0, i.e. no safety margin, is used. Without a safety margin, many more spurious obstacle candidates are created. However, extending the tracking cutoff to 20 scans allows the cinder block to be detected with no false positives. This illustrates how a trade-off in detection capability versus false positives can be made in both the intensity domain and the temporal domain.

Experiment 14 Details. A patch of bright ground pixels similar to that in Experiment 12 causes some false positives in this experiment, but to a lesser degree. A tracking cutoff of 5 results in some false positives. These false positives are eliminated if a tracking cutoff of 10 is used.

7.1.2.3 Experiments with the Environmental Cover

To test the effect of the environmental cover on the system's detection capabilities, I performed several experiments with the cover in place (see Section 4.1.2 on page 54 for a description of the cover). Vehicle position data was also used in the experiments.

Table 7-3. Tracking with Position Data, with Environmental Cover

Experiment	Number and Type of Actual Obstacles	Total candidates detected	Candidates seen at least 5/10/20 times	False Positives 5/10/20	Misses 5/10/20	Range (meters)
15	4 total: 1 cinder block, 2 lampposts, 1 car	23	3/3/3	0/0/0	1/1/1	45
16	6 total: 1 cinder block, 2 lampposts, 3 cars	64	45/31/19	0/0/0	0/0/1	70
17	7 total: 1 deer hide, 2 lampposts, 4 cars	52	31/23/16	0/0/0	0/0/0	60
18	6 total: 1 deer hide, 2 lampposts, 3 cars	28	18/15/12	0/0/0	0/1/1	45
19	5 total: 1 white-painted surface, 2 lampposts, 2 cars	39	20/16/11	0/0/0	0/0/1	46
20	4 total: 1 white-painted surface, 2 lampposts, 1 car	32	13/8/7	0/0/0	0/0/0	46

As discussed in Chapter 4 and Chapter 6, the environmental cover can cause internal reflection problems with the laser sensor. This made detection more challenging. In particular, the optical crosstalk magnified by the cover made obstacles alternately appear darker and brighter. This made obstacles undetectable during portions of an experiment, and made them more difficult to track. For this reason, some of the obstacles were not conclusively detected until significant travel distance (up to 25 meters) had occurred. Nevertheless, most of the obstacles were detected in the experiments listed in Table 7-3 with no false positives. I now discuss the errors in more detail.

Experiments 15 and 16 Details. The cinder block in both experiments 15 and 16 was not detectable with the default safety factor of 1.5. Reducing the safety factor to 1.0, i.e. eliminating the safety margin, allows the cinder block to be detected in both experiments, but results in many spurious candidates and false positives, particularly at the left edge of the

horizontal field of view. With a safety factor of 1.0, a tracking cutoff of 10 results in 2 false positives in experiment 15 and 1 in experiment 16. These false positives and most of the spurious candidates are located at the left edge of the horizontal field of view. It may be that the concentration of false positives at the edge of the HFOV are due to internal reflections off the environmental cover.

Experiments 18 and 19 Details. Both experiments 18 and 19 result in detection failures if large tracking cutoffs are used. In experiment 18, a deer hide is detected more than 5 times, but less than 10 times. In experiment 19, a white-painted surface is detected more than 10 times but less than 20. In each experiment, significant pitching motion of the vehicle causes the corresponding object to pass in and out of the sensor vertical field of view. The pitching motion does not affect the detection of the other (taller) objects in either experiment.

7.1.3 Detection Results Summary

The obstacle detection and tracking results appear promising. As with almost any obstacle detection system, performance is characterized by a trade-off between two types of detection failures: missed obstacles and false positives. All of the obstacles listed were easily detectable at distances less than 35 meters without false positives. Most of the obstacles appear to be detectable at distances of 50 meters or more (without false positives) with good parameter settings. In particular, the system was capable of detecting a cinder block at up to 60 meters using a safety factor of 1.0 and a tracking cutoff of 10.

For the system described in this thesis, the trade-off between missed obstacles and false positives can be adjusted in either the intensity domain by changing the safety factor, or in the temporal domain by adjusting the tracking cutoff. I have made no systematic attempt to find optimal values for these two parameters. In general, I used a safety factor of 1.5 and a temporal cutoff of 10 because the combination seemed to provide reasonable detection results without a significant number of false positives in almost all cases. Optimal parameter values will most likely depend on the environment and the detection needs in terms of minimum detectable obstacle height.

A larger safety factor would decrease the number of spurious detections, but might result in more misses (false negatives). A smaller safety factor might have allowed detection of the couple obstacles that were missed, but would have increased the number of false positives. A larger tracking cutoff would increase the minimum detectable obstacle height. More laser energy would help to extend the detection range and reduce the number of false positives.

None of the experiments listed above used obstacles that were significantly darker than the road surface. As mentioned in Section 6.6, dark objects are currently undetectable.

Section 7.2 details some additional experiments performed indoors with the laser sensor mounted on a cart for the purpose of evaluating range estimation performance. Although the tracking is trivialized in these experiments, the experimental results provide additional positive results for the obstacle classification method.

7.2 Range Estimation Results

This section provides additional range estimation results based on the methods outlined in Chapter 6, and provides additional discussion of the results. The first data comes from the tracking experiments that were analyzed in Section 7.1.2.

7.2.1 Outdoor Position-Tagged Data

The position-tagged data provides us with a variety of tracked obstacles for range estimation including vehicles, lampposts, and smaller obstacles. Unfortunately, it is difficult to gauge the effectiveness of the method on most of the surfaces from this data. The range estimation results for the two lampposts shown in Section 6.5 showed that the estimation techniques locked in on the correct range only after approximately 10 meters of travel. With the exception of the lampposts, most of the objects in the outdoor experiments could not be tracked for more than 10 meters because of the limited sensor field of view.

Short obstacles usually passed out of the vertical field of view of the sensor within 10 meters. Other vehicles were generally located at the periphery of the HFOV of the sensor,

and therefore passed out of the field of view as the Navlab approached them. The specular surfaces of vehicles also made them difficult to track consistently.

Section 7.2.2 examines some data I collected which avoid the limited FOV problem. But first I show some additional range estimation results for several of the lampposts tracked. The lampposts provide a nearly ideal surface for the range estimation technique. Because they are taller than the vehicle, they never disappear under the VFOV of the laser. In addition, the cylindrical surface looks the same from any azimuth, so changes in the vehicle position should not affect the laser incidence angle as it hits the lightpost. Figure 7-2 through Figure 7-4 display range estimation results for several lightposts.




Figure 7-2. Range estimation results for a lamppost. The oscillations in the estimates that appear every 6.5 meters are caused by crosstalk.

Just like the graphs in Section 6.5, all three of the graphs in Figure 7-2 through Figure 7-4 show that the range estimates begin to lock in on the correct range after approximately 10 meters of travel. By the end of each experiment, all the range estimates are correct to within a few meters. The three graphs also display an oscillation in the range estimates with a period of approximately 6.5 meters. These oscillations are caused by optical crosstalk (see Section 4.2.2 and Section 6.6.3 for an explanation of crosstalk and its effects). Crosstalk causes an unmodeled disturbance in the measured intensity values which

in turn causes the error in the range estimates. If the magnitude of the crosstalk was constant, it could be modeled and used to correct the range estimates. Even better, however, would be to remove the range channel and the AM modulation of the laser signal which causes the crosstalk.




Figure 7-3. Range estimation results for another lamppost. Again, the oscillations in the estimates that appear every 6.5 meters are caused by crosstalk.




Figure 7-4. Range estimation results for another lamppost. Again, the oscillations in the estimates that appear every 6.5 meters are caused by crosstalk. The sudden jump in the measured range is caused by the 52 meter ambiguity interval of the laser. When tracking first begins, the lamppost is approximately 67 meters

away which is equivalent to the measured value of 15 meters. Note that this range ambiguity does not fool the intensity-based measurement. Some of the estimation methods begin to oscillate about the correct range after 10 meters of travel, when the object is still 57 meters away (and measured as 5 meters away by the range channel).

One feature of the intensity-based range estimation method is that it is not subject to range ambiguity problems inherent in AM modulated lasers (see Section 4.1.1.3 for more details on AM lasers or Figure 7-4 for an example). As such, these techniques could complement systems with direct laser range information. The intensity data could be used to disambiguate the returned range values. Using the technique to disambiguate range values should be fairly easy, and should require far fewer samples than required to estimate range from intensity alone.

7.2.2 Range Estimation with an Extended Field of View

Since short obstacles quickly pass out of the sensor VFOV and since pitch variations can cause our vehicle-based laser setup to miss obstacles entirely, I decided to collect some data in a different manner to more completely test the range estimation algorithms. The laser was placed on a cart and wheeled towards the objects. A full laser image was taken every few feet. A tape measure was used to measure the cart location for each image. Then five lines which contained the object of interest were selected by hand from each image and concatenated into a single push image. Since slight yawing of the cart between images caused sudden jumps in the apparent obstacle position (by up to 60 pixels) and made automatic tracking impossible, the selected lines were also aligned horizontally by hand.

The resulting images were similar in appearance to the vehicle-based push images, and were fed through the algorithm described in Section 6.1 through Section 6.5. Although the manual horizontal alignment of the scan lines made the tracking problem trivial in this case, line-by-line obstacle classification and range estimation was performed automatically as usual. The following graphs present the results of the intensity-based range estimation from these experiments for a number of different obstacles.




Figure 7-5. Range estimation during tracking of a 2 ft. high wooden crate at an actual range of 46 meters (as measured by the laser range channel). As object tracking begins, the methods for range estimation cannot accurately estimate the object range. After approximately 7 meters of travel, the estimates begin to lock in on the correct range. After nearly 25 meters of travel, all methods correctly estimate the range to within 2 meters.




Figure 7-6. Range estimation for a cinder block. The cinder block provides a very good surface for range estimation given its uniform color and rough, diffuse surface. After only 10 meters of travel, all estimates are within 5 meters of the measured range. By the end of the experiment, all of the estimates are within 1 meter of the measured range. Relatively small crosstalk-related 6.5 meter oscillations in the range estimate are also evident here.

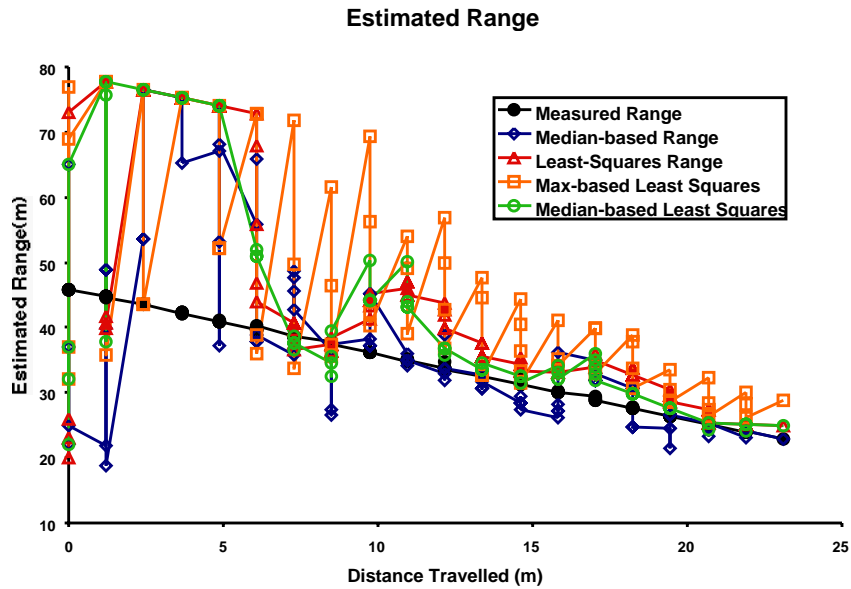


Figure 7-7. Range estimation for a deer hide. Most of the range estimates perform fairly well on this surface. Most are within 5 meters of the correct range after about 12 meters of travel, and within 2 meters at the end of the experiment. The max-based least squares technique, however, has very noisy estimates. This is likely caused by the non-uniform albedo over the surface (each of the five scanlines taken at each distance has a significantly different maximum albedo).




Figure 7-8. Range estimation for a polished metal surface with its surface normal pointed near the sensor. All of the estimation techniques underestimate the measured range at the end of the experiment. This indicates that the measured intensity rose faster than predicted by the inverse-square law. One potential reason for this is that the sensor aperture saw an increasing portion of the specular lobe as it approached the metal surface, causing a faster rise in intensity than predicted. Another potential reason for the error is optical crosstalk (as described further in Section 6.6).




Figure 7-9. Range estimation for the same polished metal surface with its surface normal pointed approximately 30 degrees away from the sensor. The least median and least-squares techniques provide good range estimates by the end of the experiment. The max-based least squares technique underestimates the range (most likely for reasons described in Section 6.5.4), while the median-based least squares overestimates the range (most likely for reasons described in Section 6.5.5). Again, some 6.5 meter crosstalk-related oscillations are evident in the graph.

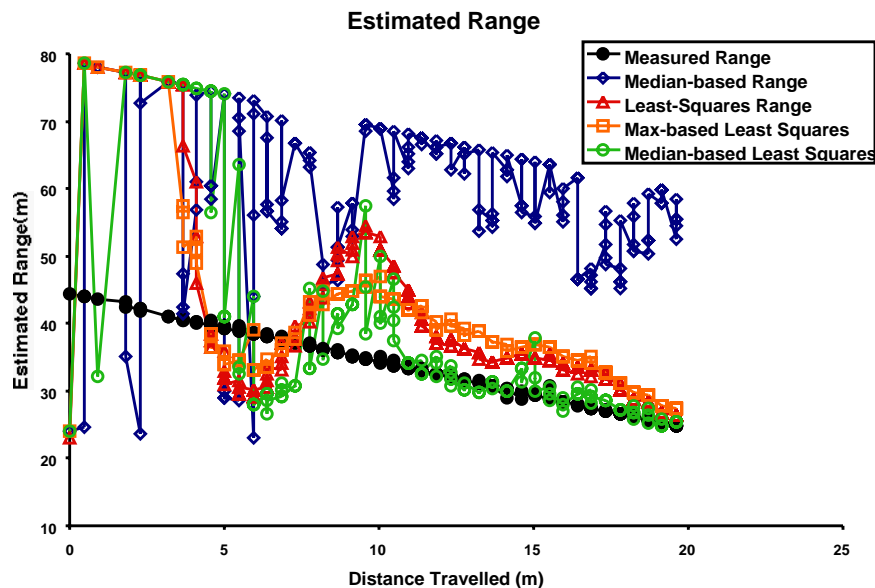


Figure 7-10. Range estimation for a striped object. The striped object is a wooden surface with a thick vertical gray stripe covering approximately 40% of the surface. The least median method performs terribly on this surface because of the heterogeneous albedo. The rest of the estimates perform reasonably, resulting in estimates within 2 meters of the measured range by the end of the experiment. Crosstalk is evident here too.

For the majority of surfaces, range estimates are reasonable (within a few meters) after 10 meters or more of travel. Surprisingly, the basic least squares estimator appears to give the best performance. Its superior performance is likely because outliers are relatively infrequent and bounded. Estimating the range to specular surfaces is difficult because small vehicle (or cart) motions can change the received energy significantly. Nonetheless, the results appear promising for most of the surfaces. Table 7-4 shows the qualitative results for the 4 range estimation techniques for the various surfaces at the end of each experiment. If the technique converged to an answer within a few meters of the correct range, it is labelled “Correct”.

Table 7-4. Qualitative Range Estimation Results for Various Obstacles

Obstacle	Least Squares	Least Median of Squares	Least Squares of Medians	Least Squares of Maxima
Lamppost	Correct	Correct	Correct	Correct
Wooden crate	Correct	Correct	Correct	Correct
Cinder block	Correct	Correct	Correct	Correct
Deer hide	Correct	Correct	Correct	Correct
Polished metal, normal oriented towards sensor	Underestimate	Underestimate	Underestimate	Underestimate
Polished metal, normal oriented 30° away from sensor	Correct	Correct	Overestimate	Underestimate
Vertically striped surface	Correct	Extreme overestimate	Correct	Correct

Looking at Table 7-4, we see that range estimation on diffuse surfaces converged correctly for most of the methods. The one exception was the vertically striped surface with the least median of squares method. This is a rather surprising result. The least median of squares method is supposed to be the most robust of the implemented statistical techniques.

To explain the poor performance of the least median of squares estimator on the striped surface shown in Figure 7-10, the raw intensity versus range distribution is graphed in Figure 7-11, along with the intensity versus distance curves generated by the various range estimators. The least squares and the least squares of medians estimators produce very sim-

ilar range and normalized reflectance estimates. The least squares of maxima method produces a much higher normalized reflectance (as expected) by fitting the maximum intensity at every range. The least squares of maxima method still produces a similar (nearly correct) range estimate to the other two methods.

The least median of squares method, however, generates a flatter curve. Since the least median of squares method only has to agree well with 50% of the data, it does not follow the general trend of the data and results in a very poor range estimate. This indicates that the least median of squares method may not work well whenever there is a large spread in the reflectance values at each range.




Figure 7-11. This plot shows all of the reflectance data collected on the striped object at the end of the experiment shown above. It also shows the curves chosen for each of range estimation method. Because the least median method can ignore almost 50% of the data, it chooses a fit that does not match the overall trend of the data.

We should expect heterogeneous surfaces to provide noisier range estimates, and the graphs seem to support this. For example, range estimates for the deer hide are noisier than for the cinder block. Heterogeneity in the vertical direction is especially problematic for range estimation because as the vehicle moves forward, the laser sweeps up the obstacle surface. The implicit assumption in the range estimation algorithm is that the top of the obstacle is similar in albedo to the bottom of the obstacle. An algorithm like the least

median of squares might be able to ignore some horizontal stripes across the surface (by ignoring this data), but an object with a smooth gradient in albedo from bottom-to-top would cause significant errors in the range estimate. Except for the least median of squares method, the methods were able to cope with horizontal heterogeneity since the laser samples an entire horizontal line across the object at each frame.

7.2.3 Range Estimation Results Summary

Range estimation results based on intensity-tracking looks promising for diffuse surfaces. All of the range estimation methods converged to within a few meters of the correct range for almost all of the surfaces.

However, it is more difficult to provide good intensity-based range estimates for specular surfaces. Specular surfaces are more difficult to track since individual specular obstacle candidates can appear or disappear fairly quickly as the laser light gets reflected towards or away from the laser receiver. In addition, the returned intensity of these specularities can vary greatly with incidence angle. This rapid variation is difficult to model or correct for without knowledge of the object shape and surface material characteristics.

One of the more surprising results of the experiments is that the least squares method performed the best of all of the range estimation methods. Theory tells us that a single outlier can cause an unbounded error in a least squares estimate. The fact that the least squares method converged to the correct range most of the time indicates that the automatic segmentation is fairly good, and that intensity outliers are infrequent and not significantly incorrect.

More heterogeneous obstacles or poor segmentation might adversely affect our range estimates. A more robust method might perform better in such cases. The least median of squares method performed badly on the striped object because it modeled an unrepresentative or biased sample of the reflectance data. Perhaps a method that tried to match 90% of the data would perform better and still allow for up to 10% outliers. Alternatively, a two-pass range estimation method might solve the problem. If the least-squares estimate is not

drastically wrong, it could be used to find outliers. These outliers could then be filtered out, and a least squares estimate could be run again on the remaining data.

CHAPTER 8 **Conclusion**

8.1 Thesis Summary and Contributions

Highway-based obstacle detection is a challenging problem. Fast highway travel speeds force long lookahead distances and fast data processing. Currently, there are no systems proven capable of reliable highway obstacle detection, particularly at night.

My thesis presents theory and implementation of a novel laser intensity-based obstacle detection system. The system is capable of detecting and tracking a large variety of objects at long distances, requires less computer processing than alternative obstacle detection systems, and potentially at a greatly reduced cost from laser ranging methods. It also offers a novel solution for target range estimation.

Unlike camera-based vision systems, this system operates best at night and is capable of detecting objects beyond the reach of headlights. Segmentation of potential obstacles is very fast and requires minimal processing. Range estimation is computationally more expensive. As such, the detection methods in this thesis would serve as an excellent complement to a stereo-based vision system. Obstacle candidates could be spotted quickly with

the laser, and registered stereo could verify a candidate and estimate its range relatively quickly.

The primary contribution of this thesis is the methodology for and demonstration of the use of laser intensity as a novel means of obstacle detection, tracking, and range estimation as described in Chapter 6. To my knowledge, laser intensity has never been used in this way before. The use of active lighting for rough range estimation could be used in a number of areas such as background separation and segmentation for vision applications. Chapter 7 shows encouraging results for the methods. This thesis makes additional contributions in several areas.

Chapter 2 provides valuable comparisons of current obstacle detection methods with comparative mathematical analysis of several vision-based methods.

Most obstacle detection systems have ignored basic system requirements and real road geometry and have thus made invalid simplifying assumptions. Chapter 3 uses the highway design manuals to show that the flat-world model is inaccurate for purposes of small obstacle detection. In contrast, it shows that a locally flat model is sufficiently accurate. This is the only such highway design-based analysis in the obstacle detection literature to my knowledge.

Chapter 4 discusses the design and analysis of a state-of-the-art laser scanner, including the operational theory of the basic rangefinding methods. It also covers many of the real-world performance issues and problems with AMCW laser scanners, and thus serves as a primer for anyone considering the design or analysis of a laser rangefinding system.

Finally, Chapter 5 examines the physics of laser reflection. In particular, it verifies the intensity versus range relationship which is then used in Chapter 6 for range estimation. In addition, it analyzes intensity noise and uses diffuse reflectance theory by Oren and Nayar[39] and specular reflectance theory by Torrance and Sparrow[48] to produce a more complete laser reflection model. This model will be useful for future researchers in laser intensity-based obstacle detection.

8.2 Future Work

As the first real system (to my knowledge) to attempt laser intensity-based obstacle detection, there are a number of areas in which the work may be extended and improved.

8.2.1 Designing a Better Sensor

To achieve significantly better obstacle detection results with the methods presented in this thesis or similar laser intensity-based techniques, an improved laser scanner design tailored to these techniques is needed.

8.2.1.1 Scanner Mechanism Modifications

The original purpose and use for our laser scanner was for cross-country vehicle mapping and navigation. For this purpose, the scanner needed to provide excellent range information with a 360 degree HFOV at ranges up to 30 meters. For the slower speeds of travel associated with cross-country navigation (compared to highway navigation), a larger VFOV with a slower frame rate was acceptable. These are poor design parameters for a scanner for highway obstacle detection.

Clearly a 360 degree HFOV is unnecessary for detecting static obstacles on highways. The only static obstacles which pose a threat to the vehicle are located in the vehicle's own or adjacent lanes. As shown in Chapter 3, a 22.7° HFOV is probably sufficient if the sensor is unsteered. With a polygonal mirror system, this might allow up to 15 lines to be collected in a single 360° rotation of the mirror, and achieve a much faster line rate without increasing the RPM of the motor. The relatively slow line rate of the current sensor has been a major factor preventing detection at highway speeds.

Although high-speed vertical mirror nodding is unnecessary, it would be desirable to have active control of the vertical (pitch) axis, to automatically compensate for pitch variations caused by the vehicle suspension. Ideally, the pitch of vehicle and sensor should match the local slope of the road. However, the suspension generally results in differences between the two. Currently, vehicle pitch variations cause missed detections and invalidate surface coverage guarantees when the scanner is operated in single scanline mode. To com-

compensate for relative pitching in the suspension, we would need an accurate inclinometer attached directly to the sensor, and another attached to an axle or the vehicle chassis below the suspension. Subtraction of sensor pitch from the chassis pitch should allow compensation for variation in the suspension. Alternatively, the sensor could be attached to the chassis without going through the suspension. However, it may be difficult to attach it in this way and position the sensor at a height where it is able to see over hill crests (approximately 1 meter high, see Section 3.2.2). Compensation for road surface curvature to maintain a consistent lookahead distance would be more difficult, but might be possible through improved road reflectance modeling.

As an alternative to active pitch control, a scanner with a small VFOV (approximately 5 degrees) would be useful. A small VFOV would tolerate small pitch variations without losing sight of the roadway or the obstacles. Our current scanner has a VFOV that is much too large for highway travel; it is impossible to simultaneously achieve adequate frame rates and vertical resolution with such a large VFOV without significantly increasing motor speeds and data throughput. There are some scanner designs that use a 1-DOF scanning motor with multiple fiber-optics to simulate a 2-DOF scanner. Such a configuration might maintain high frame rates while providing adequate vertical resolution.

The current scanner provides insufficient azimuth resolution to detect narrow objects much further than 40 or 50 meters. With the current azimuth resolution of 0.06 degree per pixel, we only get a pixel every 6 cm at 60 meters, so a foot-wide obstacle is only 5 pixels wide at 60 meters. It would be desirable to quadruple this resolution to provide double the pixels on a target at twice the range. This requires increasing the resolution of the optical encoder or increasing the gear ratio between the motor and yoke shaft. More recent versions of the scanner more than double the current azimuth resolution. Of course, increasing the azimuth resolution without a similar increase in the laser data rate will result in a decreased line rate. Currently, however, we are only running the laser at half its top data rate of 500 kHz.

8.2.1.2 Laser Modifications

Range measurement should be eliminated entirely to save on cost. This would eliminate the need for the amplitude modulation and demodulation circuitry as well as the digital signal processing. Elimination of the range channel would eliminate crosstalk problems. In particular, this would fix the problems in detection and intensity-based range estimation caused by optical crosstalk as described in Section 6.6. This should significantly improve the speed of convergence and the accuracy of intensity-based range estimates. Elimination of optical crosstalk problems would also reduce concerns about reflections caused by the environmental cover.

Another major recommendation for a new laser sensor would be to eliminate the coaxial design of the transmitter/receiver pair, and design the laser to achieve maximum overlap of the receiver and laser FOVs near the desired lookahead distance. This would solve a couple of problems. First, as discussed in Section 4.1.2 on page 54, designing an environmental cover would be made simpler since reflections off near surfaces would not be in the receiver FOV and would not interfere with the desired signal. Second, the laser power could be increased without needing to increase the sensor's dynamic range or increase the number of bits in the intensity A/D conversion. In the current design, increasing the power could saturate or cause damage to the receiver from strong reflections from near surfaces. Although the intensity resolution is 16-bit in our current sensor, surfaces beyond 10 meters only return values in the lower 10 bits. By designing the laser and receiver FOV so that surfaces at less than 5 meters range are invisible, we can increase the power or increase the gain on the avalanche photodiode without harming the sensor. One side effect of making the laser non-coaxial is that it would complicate the intensity versus range relationship. This would require additional modeling for the specific laser, but it should still be possible to estimate range from intensity tracking.

As an alternative to increasing the laser power, the sampling frequency might be reduced. Since noise is dependent on the received *energy* of the signal, it is possible to reduce noise by increasing the sampling time rather than increasing the laser power. Increasing the sampling time, of course, decreases the data rate of the sensor. Currently, a

very small fraction of the total data rate is utilized since most of the HFOV is ignored. Depending on the design parameters of a new sensor and eye safety issues involved, the energy for data rate trade-off may be useful to exploit.

Finally, in designing an improved reflectance sensor, the effects of temperature drift will need to be carefully considered. Temperature drift occurs fairly slowly, so it is unlikely to cause errors in range estimations based on intensity tracking. Significant temperature drift takes minutes or hours (see Figure 5-18 on page 87), while an obstacle will only be seen for a few seconds (at highway speeds) before the vehicle passes it. It does indicate, however, that a fixed reflectance model or threshold for the road will be suboptimal without taking temperature into account.

8.2.2 A Radically Different Sensor

The previous section discusses evolutionary changes to the laser sensor. A more revolutionary sensor might use a strobe of a specific optical frequency and a nearby CCD with a matching optical band filter. This should provide entire images similar to that produced by a scanning laser, but at much faster CCD camera rates. By eliminating optical frequencies besides that of the strobe, most ambient light should be eliminated allowing the same intensity-based approaches to be used for obstacle detection as discussed in this thesis. With 2-D image data at each frame, obstacle tracking and segmentation may become easier and more reliable since texture-based methods can be used. Range estimation should also benefit from the increased amount of data.

Such a sensor could also be used to perform road-following tasks with greater ease than existing vision based systems since illumination would be controlled and known. The shadows that typically cause problems with vision-based methods would be essentially eliminated since the lighting and viewing directions could be nearly coincident.

One problem that such a setup might have, however, is interference between multiple similar sensors. Because of the wide instantaneous field of view (IFOV) of the strobe and CCD setup, a CCD on a vehicle might receive light from the strobe on another vehicle. In contrast, this is not a concern with our current sensor because the IFOV of the laser and

receiver are so small that it is highly unlikely that the sensors on two different vehicles will be illuminating and looking at exactly the same point in space at the same time. One way to combat the interference with the strobe and CCD setup is to divide the time and optical frequency space between the sensors. By synchronizing local sensors so that they do not fire at the same time or by using different optical frequencies, interference can be avoided.

8.2.3 Software Improvements

Additional vehicle sensors and improved laser intensity hardware should enable improved results from our system. There are some software improvements that would need to be made to take advantage of improved sensors and vehicle hardware.

First, an increase in laser power and resulting decrease in sensor noise might make edge-based segmentation methods more usable and preferable to the current statistical approaches. Surfaces with very low albedo might be detectable with edge-based methods. A combination of both approaches is likely to work best. It is unclear what is the best way to combine these two approaches.

Second, the obstacle detection system needs to know where the road boundaries are located. This will help to segment objects by eliminating potentially confusing off-road reflectance data. It will also allow the system to assign road-relative coordinates to obstacles, which is useful for the vehicle control system. It might be possible to add software to find the road boundaries automatically from the laser data. Alternatively, the system could be integrated with a road-follower or GPS and map system which would provide it with lane marker locations. Since the laser system is likely to detect lane markers as potential obstacles because of their unusual brightness, it would be helpful to have the road-follower confirm these bright areas as belonging to lane markers.

Third, since the specular surfaces on cars are problematic for the laser system, the system should be integrated with a vehicle detection system (such as radar). The radar system will detect the vehicles that the laser intensity system may or may not detect. Because the laser system is tuned for static obstacles, vehicles would confuse its tracking system. By knowing the location of vehicles in the environment, the laser system could

change its tracking assumptions for obstacle candidates that are caused by reflections from vehicles, or ignore them. Since everything besides vehicles on the road should be static or nearly so, the tracking algorithms should work for everything else.

Fourth, the addition of automatic sensor pitch control should help to eliminate missed detections caused by vehicle pitching. This should enable smarter tracking algorithms. Currently, if the system spots a potential obstacle, it takes a fairly long time for the system to “expire” the obstacle if the system fails to see it again and move it to the inactive list. This long expiration time is used because vehicle pitch variation may keep the obstacle from being seen for short periods. If an obstacle candidate’s absence cannot be attributed to pitch variations, the candidate may be more quickly classified as spurious. In addition, obstacle height can be estimated from the number of lines in which the object appears (as shown in Section 3.1). This should allow the system to better analyze the threat an obstacle poses.

Finally, the range estimation algorithms will need to be adjusted to account for non-linear motion of the vehicle. Since Navlab 5 did not possess a yaw sensor, I drove in a straight line as much as possible when tracking obstacles and assumed linear travel in the range estimation algorithms. An additional complication in range estimation arises if the coaxial design of the laser is eliminated. The intensity versus range relationship will need to be studied and modeled in this case to account for the effect of differing transmitter/receiver FOV at different ranges. The new model will need to be used to in the residual formulation to obtain good range estimation.

8.2.4 Additional Testing

Given some of the shortcomings of the current laser scanner, it has been impossible to test the thesis methods in a number of areas. With an improved sensor system, additional testing needs to be performed. First, the detection system must be tested on real highways. Unfortunately, the current system could not be satisfactorily tested on highways. The line rate of the current sensor was too slow to be used at highway speeds. In addition, the environmental cover interfered greatly with the intensity data, and I could not risk damage to the sensor that might be caused by driving with the sensor uncovered at high speeds.

Impacts by small pebbles or airborne debris could seriously damage the mirror at highway speeds. The data collected so far has primarily been done in parking lots where it was safe to drive slowly without the environmental cover.

With additional laser power, the sensor must also be tested during the daytime. The current sensor could not be effectively used during the day since ambient photon noise was often stronger than the signal at long distances. All data displayed in this thesis was gathered at night or indoors where ambient photon noise was minimal.

8.2.5 Active Lighting in Other Domains

Computer vision with uncontrolled lighting poses a difficult problem. The laser system simplifies the vision problem significantly by controlling the lighting. In particular, the coincident lighting and viewing directions make the intensity-based obstacle detection problem feasible.

However, a laser is not necessary to use the ideas presented in this thesis. The laser provided a convenient testbed for the work because of its 16-bit intensity channel. Similar active lighting methods might work well for other robots. For example, a sentry robot that operates at nighttime in a dark building could use a light and CCD to distinguish walls, floors, and openings. If operation in a lit building was required, a laser or a frequency-specific strobe and matching CCD would have to be used. Similar range estimation procedures to the ones described in this thesis could be used to estimate distance to walls.

Sonar sensors are popular for indoor use, but have the problem that walls often cause specular reflections with ultrasound. This can cause ultrasonic sensors to report there is no wall in a certain direction even when there is. An individual ultrasonic measurement is therefore unreliable. Because of its shorter wavelength, light has no such problem.

Since painted walls may have nearly the same albedo over large regions, rough range estimation or planar surface extraction and mapping might be performed from an intensity image using optimization on a surface orientation and distance parameter, given a rough albedo estimate from previous mapping with coincident lighting and viewing.

Vision problems rarely have simple solutions. However, there is a wealth of information provided by visual data if the right models and methods are used to exploit it. This thesis provides a computationally and conceptually simple method for obstacle detection by doing just that.

APPENDIX A **A Simplified
Torrance-Sparrow
Model**

The Torrance-Sparrow reflectance model is a bidirectional reflectance model for roughened specular surfaces. As explained in Chapter 5, bidirectional reflectance is parameterized by both the light source direction and the viewing direction. For laser-based vision, however, we can use a simplified model where the light source and viewing directions are always coincident. In this case, the reflectance can be parameterized by just θ , the angle between the macro surface normal and the laser beam.

A derivation for the laser vision case is presented here since some of the formulas in the original Torrance and Sparrow paper are degenerate and unusable for this special case. Sato and Ikeuchi[46] presented a similar model for the coincident lighting and viewing case, and our simplified Torrance and Sparrow model is nearly identical to that of Sato and Ikeuchi. However, my model is also valid at high incidence angles whereas their model is only valid for incidence angles less than 45 degrees.

The Torrance and Sparrow model was designed to explain reflection from rough surfaces. The model assumes that a surface is made of long, symmetrical V-shaped micro-cavities. “Long” means that cavity length is much larger than the cavity width. The model is valid when the wavelength of the light is much smaller than the surface roughness, and when the spot size of the laser is much larger than a facet. The orientations of the cavity major axis are assumed to be random, and the cavity facet slopes are assumed to follow a normal distribution with zero mean. The model assumes a Lambertian diffuse component plus a specular component based on mirror-like facets. Specular reflections from facets obey the rule $\theta_r = \theta_i$.

The facet slope distribution is:

$$P(\alpha) = c e^{\frac{-\alpha^2}{2\sigma_\alpha^2}}$$

Equation A-1. The facet slope distribution. c is a constant, and the facet slope α has a mean value of zero and standard deviation σ_α .

The geometry for the general Torrance and Sparrow model is described by three angles. Direction of incident light is specified by the zenith angle θ_i measured from the surface normal. Reflected light direction is specified by both the zenith angle θ_r measured from the surface normal, and the azimuth angle ϕ measured from the plane of incidence. Only facets having normal vectors within the solid angle $d\omega$ can reflect incoming light specularly

Substituting, we have:

$$dL_r = \frac{G(\theta)L_i d\omega_i fP(\alpha)d\omega dA \cos\gamma}{d\omega_r dA \cos\theta_r}$$

The solid angle $d\omega$ which contains the surface normals of interest can be related to the solid angle of reflection $d\omega_r$:

$$d\omega_r = 4d\omega \cos\gamma$$

For the laser vision case, we have the following diagram:




Figure A-2. A simplified reflection diagram for the Torrance and Sparrow model for the laser vision case.

We can substitute $\theta_i = \theta_r = \alpha = \theta$, $\gamma = 0$, and $\phi = 0$. We can also consider the fresnel coefficient F a constant for an individual surface. Simplifying, we have:

$$dL_r = k \frac{G(\theta)P(\theta)}{\cos\theta}$$

Equation A-2. Expression for the surface radiance for laser reflectance. k is a normalizing constant.

For coincident lighting and viewing rays, specular reflections are only received by the laser detector if the ray is normal to the facet surface. In this case, the angle of incidence to the macro-surface normal is equal to the facet inclination angle.

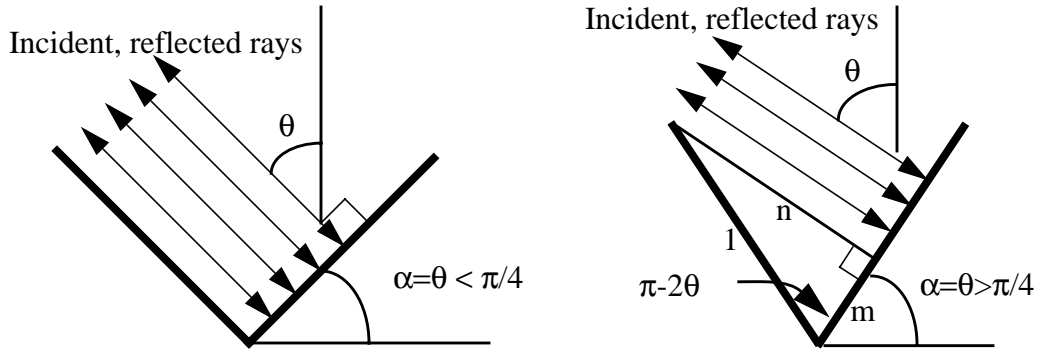


Figure A-3. Reflection diagram for specular V-groove cavities. For the mirror-like facet, light is reflected to the laser receiver if and only if the facet slope α equals the incidence angle θ . For cavity slopes of 45 degrees or less, there is no facet masking or shadowing for coincident rays. At slopes greater than 45 degrees, however, a portion of the facet is shadowed by the opposite facet. This shadowed portion is of width m . Without loss of generality, we assign each facet a width of one.

The geometric attenuation factor $G(\theta)$ is the fraction of the facet surface that contributes to the reflected flux. From Figure A-3, G can be defined as: $G(\theta) = 1 - m$.

G is equal to one for incident angles less than 45 degrees since the light rays are unblocked by the opposite facet, but decreases to zero as the incident angle approaches 90 degrees (see Figure A-3). We can write the following equations to describe the reflection triangle:

$$\cos(\pi - 2\theta) = -\cos(2\theta) = m$$

This gives a simple two-case expression for G :

$$G(\theta) = 1, 0 \leq \theta \leq \frac{\pi}{4}$$

$$G(\theta) = 1 + \cos(2\theta) = 2 \cos^2 \theta, \frac{\pi}{4} \leq \theta \leq \frac{\pi}{2}$$

Equation A-3. The geometric attenuation factor, G is the fraction of the facet surface that contributes to the reflected flux. G is a function of the incidence angle θ .

If we now add the Lambertian diffuse component, we have:

$$I(\theta) = k\rho \left(\frac{(1-C)}{\cos\theta} e^{\frac{-\theta^2}{2\sigma_\alpha^2}} + C \cos\theta \right), \text{ for } 0 \leq \theta \leq \frac{\pi}{4}$$

$$I(\theta) = k\rho \cos\theta \left(2(1-C) e^{\frac{-\theta^2}{2\sigma_\alpha^2}} + C \right), \text{ for } \frac{\pi}{4} < \theta \leq \frac{\pi}{2}$$

Equation A-4. A simplified Torrance and Sparrow model including a Lambertian diffuse component. θ is the angle of incidence, ρ is the albedo, k is a normalizing constant, and C is the percentage of energy reflected diffusely. σ_α is the standard deviation of facet slope.

Equation A-4 gives the simplified Torrance and Sparrow model for the special case for the laser, and depends on two surface parameters: the standard deviation of the facet slope σ_α , and C , the percentage of light reflected diffusely.

APPENDIX B

A Kalman Filter for Obstacle Position Estimation

To enable better tracking and localization of an obstacle, it is useful to have a target range estimate. As described earlier in the thesis, laser rangefinders are expensive, and we would like to perform obstacle detection entirely without active range sensing. Without direct range information, however, we can still produce reasonable range estimates based on the laser intensity. Since laser intensity, I , can be related to range by $I \propto 1/R^2$, where R is the target range, we can get estimate range to a target if we can track it over many successive frames. An extended kalman filter was created to take advantage of this relationship.

When we first see a potential obstacle, we can estimate its range based on the inclination of our sensor ray. This estimate may have a large uncertainty associated with it because small variations in the sensor inclination create large variations in the range. Range estimates based on the intensity can improve upon this rough estimate.

We only attempt to track the obstacle in vehicle-relative coordinates, specifically target range and target azimuth. To use the intensity information, we also track the normalized reflectance of the target, λ . Our state vector is then $X=(\lambda,R,\varphi)^T$ and its covariance matrix is called C . A subscript on the state vector or covariance matrix indicates a given time instant.

The first step of the kalman filter considers vehicle motion between scans. Given the motion estimate of the vehicle, $Y = (s,\gamma)^T$, where s is the forward vehicle motion (along the vehicle axis), and γ is the change in vehicle yaw (assumed at the end of forward motion). We illustrate the scene




Figure B-1. We update our estimate of the obstacle position using vehicle egomotion parameters. We model the vehicle motion as a forward motion of length s , followed by a change in yaw, γ . The vehicle-relative position of the obstacle is in polar coordinates, where R is the range to the target, and φ is the azimuth.

We have the state update equations:

$$\begin{aligned}\lambda_{i+1} &= \lambda_i \\ \varphi_{i+1} &= \text{atan}\left(\frac{x}{y-s}\right) - \gamma \\ R_{i+1} &= \sqrt{(y-s)^2 + x^2} \\ x &= R_i \sin \varphi_i \quad y = R_i \cos \varphi_i\end{aligned}$$

To update the covariance matrix, we need the two Jacobians of the new state vector with respect to the old state and the motion vector:

$$J_1 = \frac{\partial X_{i+1}}{\partial X_i} = \begin{bmatrix} \frac{\partial \lambda_{i+1}}{\partial \lambda_i} & \frac{\partial \lambda_{i+1}}{\partial R_i} & \frac{\partial \lambda_{i+1}}{\partial \varphi_i} \\ \frac{\partial R_{i+1}}{\partial \lambda_i} & \frac{\partial R_{i+1}}{\partial R_i} & \frac{\partial R_{i+1}}{\partial \varphi_i} \\ \frac{\partial \varphi_{i+1}}{\partial \lambda_i} & \frac{\partial \varphi_{i+1}}{\partial R_i} & \frac{\partial \varphi_{i+1}}{\partial \varphi_i} \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{R_i - s \cos \varphi_i}{\sqrt{k}} & \frac{R_i s \sin \varphi_i}{\sqrt{k}} \\ 0 & \frac{-s \sin \varphi_i}{k} & \frac{R_i (R_i - s \cos \varphi_i)}{k} \end{bmatrix}$$

$$J_2 = \frac{\partial X_{i+1}}{\partial Y} = \begin{bmatrix} \frac{\partial \lambda_{i+1}}{\partial s} & \frac{\partial \lambda_{i+1}}{\partial \gamma} \\ \frac{\partial R_{i+1}}{\partial s} & \frac{\partial R_{i+1}}{\partial \gamma} \\ \frac{\partial \varphi_{i+1}}{\partial s} & \frac{\partial \varphi_{i+1}}{\partial \gamma} \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ \frac{s - R_i \cos \varphi_i}{\sqrt{k}} & 0 \\ \frac{R_i \sin \varphi_i}{k} & -1 \end{bmatrix}$$

$$\text{where } k = R_i^2 + s^2 - 2sR_i \cos \varphi_i = R_{i+1}^2$$

The kalman filter compound equation then tells us:

$$C_{i+1} = J_1 C_i J_1^T + J_2 C_Y J_2^T$$

$$\text{where } C_Y = \begin{bmatrix} \sigma_s^2 & 0 \\ 0 & \sigma_\gamma^2 \end{bmatrix}$$

After the compound step, we have a new estimate for the vehicle-relative location of the obstacle with compounded uncertainty. The next step uses the laser scan to produce a new obstacle state estimate with reduced uncertainty.

The laser scan provides us a state vector $Z = (I \ \varphi)^T$ where I is average intensity of the segmented obstacle, and φ is the obstacle azimuth. Our state estimate, of course, is still $X = (\lambda, R, \varphi)^T$. We have a function f that relates our measurement variables to the state variables:

$$f = \begin{bmatrix} f_1 \\ f_2 \end{bmatrix} = \begin{bmatrix} I - \frac{\lambda_i}{R_i^2} \\ \varphi - \varphi_i \end{bmatrix} = 0$$

The Kalman gain, K , is given by:

$$K = C_i H^T (W + H C_i H^T)^{-1}$$

$$H = \frac{\partial f}{\partial X_i} = \begin{bmatrix} -\frac{1}{R_i^2} & \frac{2\lambda}{R_i^3} & 0 \\ 0 & 0 & -1 \end{bmatrix}$$

$$W = \left(\frac{\partial f}{\partial Z} \right) C_Z \left(\frac{\partial f}{\partial Z} \right)^T = C_Z = \begin{bmatrix} \sigma_I^2 & 0 \\ 0 & \sigma_\varphi^2 \end{bmatrix}$$

The new state vector and its covariance matrix are given by:

$$X_{i+1} = X_i - K f(X_i, Z)$$

$$C_{i+1}^{-1} = C_i^{-1} + H^T W^{-1} H$$

The variance values in C_Z must be estimated. Dividing the standard deviation of intensity as described in Chapter 5 by the number of pixels on the object gives us the value for σ_I . We estimate the standard deviation of the azimuth error to be approximately 1 pixel or $2\pi/6000 = 0.001$ radians, assuming that points near the edges of objects may be misclassified occasionally.

References

- [1] American Association of State Highway and Transportation Officials (AASHTO), *A Policy on Geometric Design of Highways and Streets*. Washington, D.C., 1984.
- [2] Ancona, N. A Fast Obstacle Detection Method based on Optical Flow. *Proceedings of the European Conference on Computer Vision (ECCV '92)*, 1992.
- [3] Bazin, G. and B. Journet. A new laser range-finder based on FMCW-like method. *Proceedings of the IEEE Instrumentation and Measurement Technology Conference*, 1996.
- [4] Bishop, J. R. Intelligent Travel: The Automated Highway System. *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS-3)*, 1993.
- [5] Bohrer, S., M. Brauckmann, and W. von Seelen. Visual Obstacle Detection by a Geometrically Simplified Optical Flow Approach. *10th European Conference on Artificial Intelligence Proceedings (ECAI 92)*, 1992.
- [6] Bohrer, S., T. Zielke, and V. Freiburg. An Integrated Obstacle Detection Framework for Intelligent Cruise Control on Motorways. *Proceedings of the Intelligent Vehicles '95 Symposium*, 1995.

-
- [7] Bruyelle, J.-L. and J.-G. Postaire. Direct Range Measurement by Linear Stereovision for Real-Time Obstacle Detection in Road Traffic. *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS-3)*, 1993.
- [8] Cornell, S., J. Porrill, J.E.W. Mayhew. Ground Plane Obstacle Detection Under Variable Camera Geometry Using a Predictive Stereo Matcher. *Proceedings of the British Machine Vision Conference (BMVC '92)*, 1992.
- [9] Dickmanns, E. Performance Improvements for Autonomous Road Vehicles. *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS-4)*, 1995.
- [10] Eberle, K. Reflectance Processing. Erim Staff Report, *Range and Reflectance Processing Workshop Proceedings*, Warren, MI, December 1987.
- [11] Enkelmann, W. Obstacle Detection by Evaluation of Optical Flow Fields from Image Sequences. *Image and Vision Computing (UK)* vol. 9, no. 3, June 1991.
- [12] Frölich, C., M. Mettenleiter, F. Haertl. Imaging Laser Radar for High-Speed Monitoring of the Environment. *Proceedings of the SPIE Conference on Intelligent Transportation Systems*, 1997.
- [13] Guinand, Y., et. al. Low Cost Sensors for Collision Avoidance Applications. *Proceedings of the Intelligent Vehicles '95 Symposium*, Sep. 25-26, 1995., pp. 478-82.
- [14] Hancock, J., M. Hebert, and C. Thorpe. "Laser Intensity-Based Obstacle Detection." *Proceedings of the IEEE Conference on Intelligent Robots and Systems (IROS)*, 1998.
- [15] Hancock, J., E. Hoffman, R. Sullivan, D. Ingimarson, D. Langer, M. Hebert. "High-performance laser range scanner." *Proceedings of the SPIE Conference on Intelligent Transportation Systems*, 1997.
- [16] Hancock, J. *High-Speed Obstacle Detection for Automated Highway Applications*. Thesis proposal. Carnegie Mellon Technical Report, CMU-RI-TR-97-17, May 1997.

-
- [17] Hartley, R. and R. Gupta. Computing Matched-epipolar Projections. *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR '93)*, 1993.
- [18] Hebert, M. SMARTY: Point-Based Range Processing for Autonomous Driving. *Intelligent Unmanned Ground Vehicles: Autonomous Navigation Research at Carnegie Mellon*. Kluwer Academic Publishers, Boston, 1997.
- [19] 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.
- [20] Heisele, B. and W. Ritter. Obstacle Detection Based on Color Blob Flow. *Proceedings of the Intelligent Vehicles '95 Symposium*, 1995.
- [21] Jin, Y., et. al. Dynamic Obstacle-detecting System for Railway Surroundings Using a Highly Accurate Laser-sectioning Method. *Proceedings of International Conference on Industrial Electronics, Control and Instrumentation (IECON) '91*, 1991.
- [22] L. Kaminski, J. Allen, I. Masaki, G. Lemus. A Sub-Pixel Stereo Vision System for Cost-Effective Intelligent Vehicle Applications. *Proceedings of the Intelligent Vehicles '95 Symposium*, 1995.
- [23] Kanade, T. and M. Okutomi. A Stereo Matching Algorithm with an Adaptive Window: Theory and Experiment. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 16, No. 9, Sept. 1994.
- [24] Kelly, A. An Intelligent Predictive Control Approach to the High-Speed Cross-Country Autonomous Navigation Problem. Ph.D. thesis. *Carnegie Mellon Technical Report*, CMU-RI-95-33, 1995.
- [25] Koller, D., T. Luong, and J. Malik. Binocular Stereopsis and Lane Marker Flow for Vehicle Navigation: Lateral and Longitudinal Control. *University of California, Berkeley Technical Report UCB/CSD 94-804*, 1994.
- [26] Kress, G. and H. Almula. *Sensorimotor Requirements for Teleoperation*. FMC Corporation Report R-6279, December 1988.
- [27] Krotkov, E., R. Hoffman. Terrain Mapping for a Walking Planetary Rover. *IEEE Transactions on Robotics and Automation*, Vol. 10, No. 6, Dec. 1994.

-
- [28] Kruger, W., W. Enkelmann, and S. Rossle. Real-Time Estimation and Tracking of Optical Flow Vectors for Obstacle Detection. *Proceedings of the Intelligent Vehicles '95 Symposium*, 1995.
- [29] Kweon, I., R. Hoffman, and E. Krotkov. *Experimental Characterization of the Perceptron Laser Rangefinder*. Carnegie Mellon Technical Report, CMU-RI-TR-91-1, Jan. 1991.
- [30] Langer, D. *An Integrated MMW Radar System for Outdoor Navigation*. Ph.D. Thesis. Carnegie Mellon Technical Report, CMU-RI-97-03, 1997.
- [31] Langer, D. Personal conversation, Jul. 29, 1998.
- [32] Langer, D. J. Rosenblatt, M. Hebert. An Integrated System for Autonomous Off-Road Navigation. *Intelligent Unmanned Ground Vehicles: Autonomous Navigation Research at Carnegie Mellon*. Kluwer Academic Publishers, Boston, 1997.
- [33] Matthies, L. Stereo Vision for Planetary Rovers: Stochastic Modeling to Near Real-Time Implementation. *International Journal of Computer Vision*, 8:1, pp. 71-91, 1992.
- [34] 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.
- [35] Matthies, L., A. Kelly, T. Litwin, and G. Tharp. Obstacle Detection for Unmanned Ground Vehicles: A Progress Report. *Proceedings of the Intelligent Vehicles '95 Symposium*, 1995.
- [36] Mettenleiter, M. Personal e-mail. November 24, 1998.
- [37] 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.
- [38] Oda, K. Calibration Method for Multi-Camera Stereo Head for NavLab II. Internal CMU Document, 1996.

-
- [39] Oren, M. and S. Nayar. Generalization of the Lambertian Model and Implications for Machine Vision. *International Journal of Computer Vision*, Vol. 14, pp. 227-251, 1995.
- [40] Pont, W. Basic Characteristics of the 3-D Sensor. Erim Staff Report, Range and Reflectance Processing Workshop, Warren, MI, December 1987.
- [41] Robert, L., M. Buffa, and M. Hebert. Weakly-Calibrated Stereo Perception for Rover Navigation. *Proceedings of the International Conference on Computer Vision (ICCV)*, 1995.
- [42] Robert, L. and M. Hebert. Deriving Orientation Cues from Stereo Images. *Proceedings of the European Conference on Computer Vision (ECCV '94)*, pp. 377-88, 1994.
- [43] Ross, B. A Practical Stereo Vision System. *Proceedings of Computer Vision and Pattern Recognition (CVPR)*, 1993.
- [44] Rousseeuw, P. Least Median of Squares Regression. *Journal of the American Statistical Association*, 79: 388, December, 1984.
- [45] Ruichek, Y. and J.-G. Postaire. Real-Time Neural Vision for Obstacle Detection Using Linear Cameras. *Proceedings of the Intelligent Vehicles '95 Symposium*, 1995.
- [46] Sato, Y. and K. Ikeuchi. Reflectance Analysis for 3D Computer Graphics Model Generation. *Graphical Models and Image Processing*, vol. 58, no. 5, 1996.
- [47] Sukthankar, R., J. Hancock, D. Pomerleau, and C. Thorpe. A Simulation and Design System for Tactical Driving Algorithms. *Proceedings of AI, Simulation and Planning in High Autonomy Systems (AISP '96)*, 1996.
- [48] Torrance, K. and E. Sparrow. Theory for Off-Specular Reflection from Roughened Surfaces. *Journal of the Optical Society of America*, Vol 57, No. 9, 1967.
- [49] Wan, Y., F. Cabestaing, J.-C. Burie. A New Edge Detector for Obstacle Detection with a Linear Stereo Vision System. *Proceedings of the Intelligent Vehicles '95 Symposium*, 1995.

- [50] Williamson, T. A High-Performance Stereo Vision System for Obstacle Detection. Ph.D. Thesis. Carnegie Mellon Technical Report, CMU-RI-TR-98-24, September, 1998.

- [51] Young, G., et. al. Obstacle Detection for a Vehicle Using Optical Flow. *Proceedings of the Intelligent Vehicles '92 Symposium*, 1992, pp. 185-190.