High-Speed Obstacle Detection for Automated Highway Applications

Thesis Proposal

John A. Hancock

CMU-RI-TR-97-17

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

May 1997

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, and a DARPA/TACOM grant, “Technology Enhancements for UGVs”, DAAE07-96-C-X075.

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

Abstract

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. Although there are a number of methods that can successfully detect moving vehicles, the more difficult problem of finding small, static road debris such as tires or crates remains unsolved. Systems such as the Automated Highway System (AHS) which demand high levels of safety are not feasible unless these critical problems are addressed. Although the problem of detecting static obstacles has been tackled in both the cross-country and indoor mobile robot navigation literature, these systems have operated at low speeds (5 - 10 mph or less) and short range.

This thesis will improve on the current state-of-the-art, by demonstrating how small static road debris can be safely detected at long distances and high speeds. In particular, it will focus on using two sensor modalities: laser reflectance and stereo vision. Laser reflectance, to our knowledge, has not been used for obstacle detection before. The thesis will show that reliable detection can be achieved by using the right methods (sensitive enough) and the right models (no more complicated than necessary) for both road and sensor.

We will develop two detection systems, laser and stereo-based, which can detect 20 cm high obstacles at 60 meters. The first system uses laser intensity to provide a more direct means of measuring surface orientation than traditional laser range-based processing: vertical obstacles should provide stronger laser returns than the horizontal road. The second system proposed is a predictive, model-based stereo method. Accurate modeling of the road and CCD sensor will enable obstacle detection without expensive 3-D reconstruction.

Table of Contents

1. Introduction

Obstacle detection is one of the key problems in computer vision and mobile robotics. Because of its necessity for mobile robot platforms, there has been extensive work on obstacle detection in a number of domains. 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 collision avoidance problem, much of it driven by programs such as AHS (Automated Highway System) [4] or PROMETHEUS which seek to revolutionize automobile transportation, providing consumers with a combination of “smart” cars and smart roads.

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],[13],[16],[18],[29],[31] 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],[16],[18], but the more difficult problem of finding small, static road debris such as tires or crates remains unsolved. Systems such as AHS which demand high levels of safety are not feasible unless these critical problems are addressed [4]. Although the problem of detecting static obstacles has been tackled in both the cross-country [15],[21],[26] and indoor mobile robot navigation literature [2],[5],[8] these systems have operated at low speeds (5-10 mph) 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. My thesis will demonstrate how small static road debris can be safely detected at long distances and high speeds. In particular, it will focus on using two sensor modalities: laser reflectance and stereo vision. Laser reflectance, to our knowledge, has not been used for obstacle detection before. While it is fruitless to account for every possible complicating factor, this thesis will show that more reliable detection can be achieved by using the right methods (sensitive enough) and the right models (no more complicated than necessary) for both road and sensor. Currently, a locally planar road model appears to be adequate (See section 4. on page 6). Sensor models are discussed in Sections 5 and 6.

We will develop two detection systems, laser and stereo-based, which can detect 20 cm high obstacles at 60 meters. Additionally, we will produce the necessary sensor signal models and geometric road models to achieve this goal. Sensor models will include noise models for our CCD cameras and signal processing models for the laser devices. Since knowing the failure modes of any system is important, we will perform sensitivity analyses of the algorithms to changing geometric parameters. Some sensitivity analyses may be done with the aid of our highway simulator, SHIVA [30], since it can produce images with known camera orientations and road geometry. We plan to integrate the two detection systems to provide greater reliability, although this may be constrained by vehicle hardware limitations (power systems, etc.).

2. Basic System Safety Requirements

There are a number of system-level requirements and basic difficulties that are associated with high-speed obstacle detection regardless of the sensor system used. We'll address these first before we examine stereo or laser in particular.

Kelly, in his Ph.D. thesis, codified the major safety requirements for a cross-country autonomous navigation system[15]. 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 $\mathbf{v}_0$ is the initial velocity of the vehicle, $\mathbf{a}$ is the braking deceleration ($\mathbf{a} < 0$) and $\mathbf{t}_{\text{sense}}$, $\mathbf{t}_{\text{process}}$, and $\mathbf{t}_{\text{brake}}$ are the latencies associated with sensor acquisition, computer processing, and braking reaction.

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

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

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:

$$\rho_{cyc} = \frac{vT_{cyc}}{\Delta R}$$

where $\mathbf{T}_{\text{cyc}}$ is the cycle time and $\Delta\mathbf{R}$ is the projection of the pixels examined by the algorithm in a single cycle on the ground plane (see Figure 1). Equivalently, $1/\rho_{\text{cyc}}$ 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. $\Delta\mathbf{R}$ 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 constraint (see Figure 2).

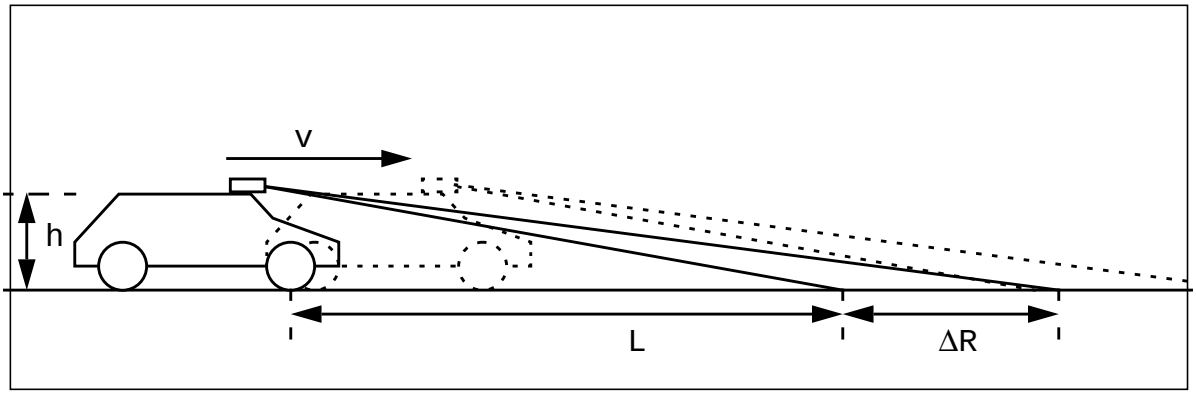

Figure 1. The vehicle must look ahead a distance L so that it can react in time. It must also examine enough terrain (ΔR) at each time step to guarantee coverage.

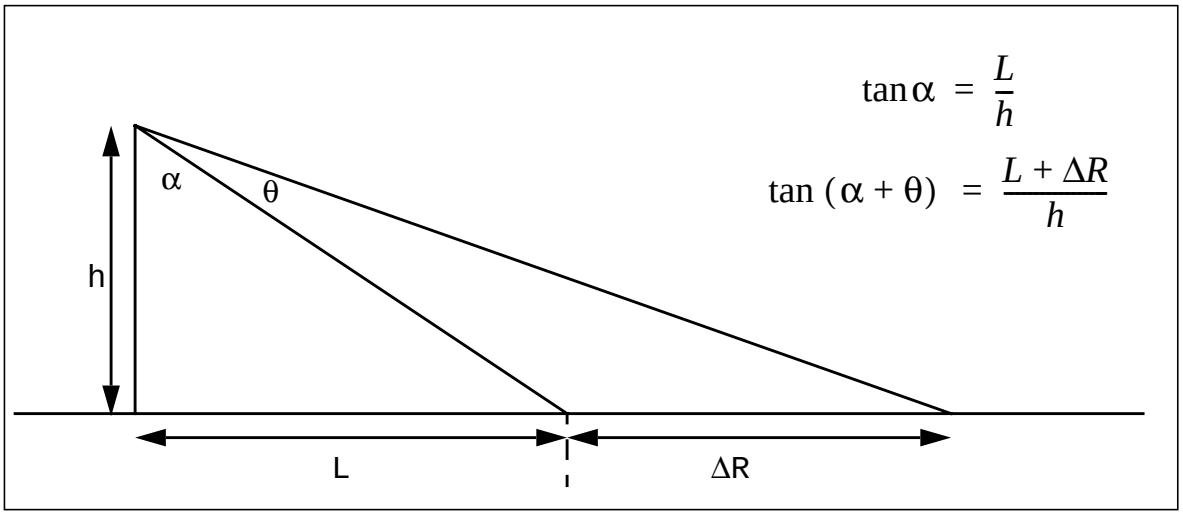

Figure 2. We can translate the throughput requirement on ΔR into a requirement on the vertical field-of-view (VFOV), shown here as θ.

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

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

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

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, based on the Nyquist sampling theorem, states that at least two pixels must fall on a

vertical obstacle for it to be detected. Throughout the paper, 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 the sensor (see Figure 3) with the following equation:

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

In general, however, this greatly underestimates the necessary vertical angular resolution since oversampling is desirable for improved reliability of the detection system. As we explain later (See section 5. on page 8), this acuity constraint is less important for the laser reflectance sensor.

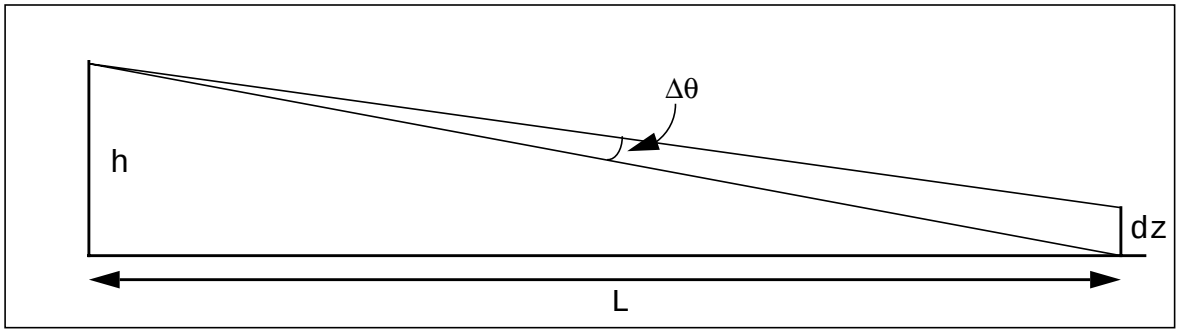

Figure 3. The vertical height a single measurement subtends, dz, is related to the vertical angular resolution of the sensor, Δθ.

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 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

Table 1: Typical Values for Highway Obstacle Detection

Variable	Meaning	Value	Given (G) / Calculated (C)
T_{cyc}	Cycle time	0.3 s	G
ΔR	ground range	8.01 m	C
h	height of sensor	1 m	G
VFOV	vertical field of view	0.0017 rad (~0.1^o)	C
$\Delta\theta$	sensor vertical angular resolution	0.0015 rad (~0.1^o)	C

We should 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. On-road versus Off-road

Although there have been good results published for vehicle detection, there have been no satisfactory solutions so far to the problem of small static obstacle detection at highway speeds. Moreover, the few results reported for static obstacle detection have generally been 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 more excusable since rough terrain is difficult to describe in an analytic fashion and an obstacle may be less well-defined in cross-country applications. There has been some success in cross-country obstacle detection, however, so it is worth briefly examining the problem.

Typical speeds for cross-country applications are on the order of 5 to 10 mph. Typical highway speeds, however, are many times this. Since stopping distance is proportional to the square of 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 sensor 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 sensor field-of-views. 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 (locally, at least). Given our previous analysis, it should be quite clear that we need to use this fact in order to win in this scenario. Although cross-country terrain processing methods need to build explicit maps (since obstacle density may be high), estimate heights of objects, and check for rollover conditions, etc., these steps are unnecessary for roadway navigation. We only need to find areas which violate the planarity constraint and avoid them.

4. Road Geometry Considerations

Before we start using the planarity constraint to help solve our obstacle detection problems, we must examine how close a highway is to being planar.

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 a minimum 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 mph, the specified **L** is 580 feet. We can calculate the minimum radius of vertical curvature of the road using the following diagram and equations.

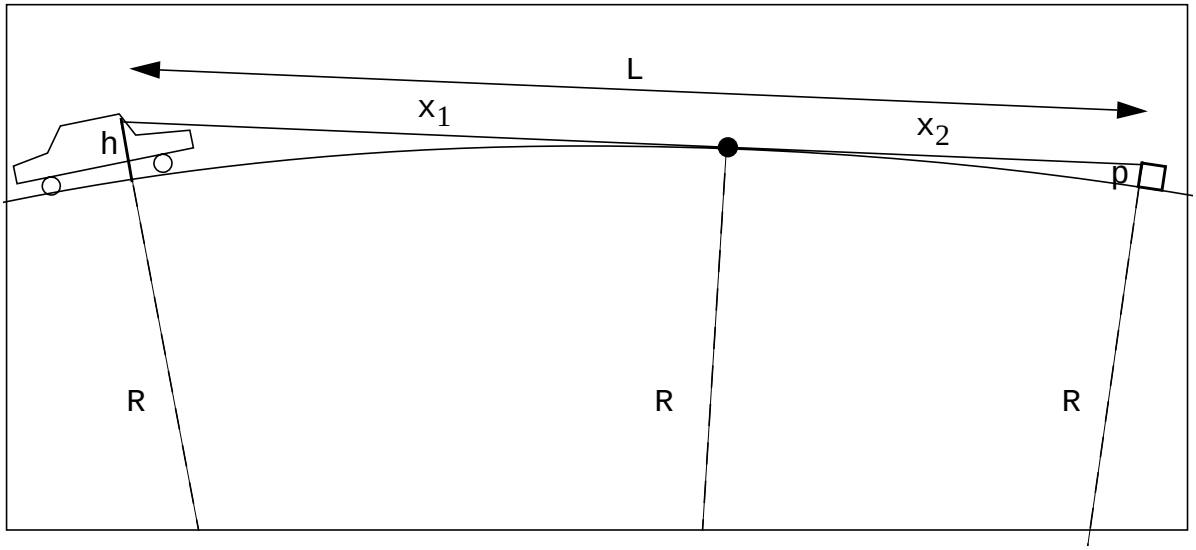

Figure 4. Highway design specifications state that a driver at height h = 3.5 feet above the road surface must be able to see an object of height p = 0.5 feet on the road at a distance L. L is dependent on the designed highway speed. Given L, we can calculate the radius of vertical curvature R.

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

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

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

have:

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

We now calculate how much the road deviates from planar over our lookahead distance, **L**. 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)?

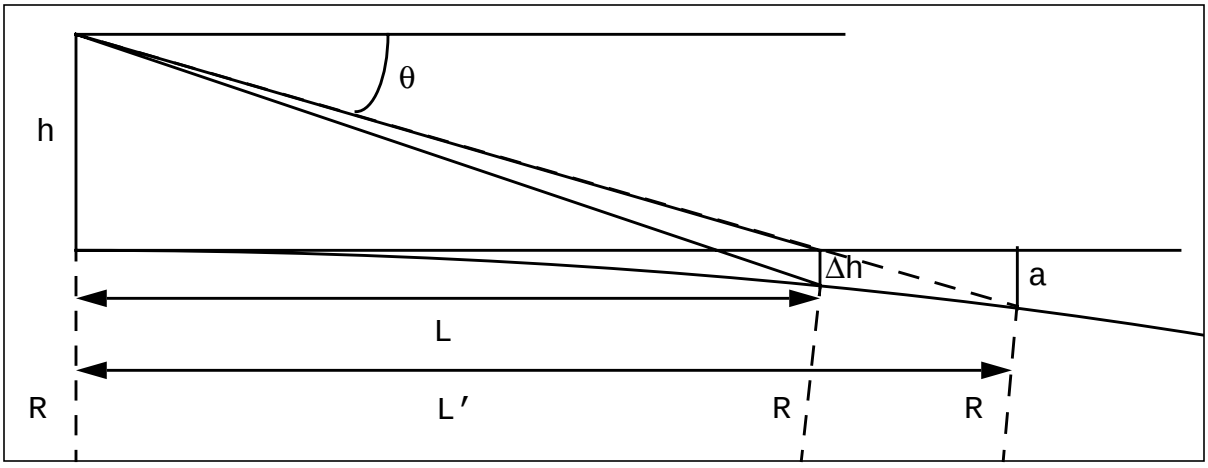

Figure 5. Our sensor is normally inclined at angle θ to intersect a planar road at distance L. However, because the road has a finite radius of curvature, R, the ray from the sensor intersects at a distance L′, and at a depth of a below the assumed surface. Δh is the depth of the road below the assumed surface at the nominal lookahead distance L.

$$\text{Let} \quad L = 60\text{m}$$

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

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

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

Substituting in for **a**, we have:

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

Coupling vehicle and sensor pitch variations with this curvature effect, we see we can not know a priori at what depth a sensor ray should intersect the road. Clearly methods that attempt to detect obstacles based on absolute depth are unlikely to work. Pitch variations in

the vehicle and sensor might force linear sensor-based methods to have cycling rates faster than the rate of pitch variation to guarantee coverage, depending on pitch variation amplitude.

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. Allowable changes in road grade for "sag" vertical curves may be nearly twice that for vertical crest curves [1]. Thus, flat world models simply can not be relied upon for small obstacle detection on many highways.

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 error x, where:

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

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.

5. Laser Reflectance

Traditional methods of laser range processing for off-road obstacle detection involve transforming the range image into an elevation map. Since the elevation map is generally sparse and noisy, smoothing is generally necessary. Then a terrain-typing algorithm is run to classify areas of the elevation map as navigable or unnavigable. Besides being a fairly complicated process, this method has another drawback in that elevation maps may not represent obstacles well, since they cannot represent vertical planes because of the discrete grid size [3]. On the other hand, laser reflectance ought to provide us with a more direct means of finding obstacles or vertical surfaces. At the long lookahead distances and grazing angles typical of high-speed travel, horizontal surfaces should provide very weak (or nonexistent) laser returns. Vertical surfaces, however, should result in stronger signals.

Preliminary tests with the ERIM laser rangefinder have shown that laser reflectance can detect small obstacles at ranges of around 20 meters using simple processing (see Figure 6 to Figure 8). The image in Figure 8 was obtained without explicit sensor modeling. A quadratic curve was fit to all reflectance values of road pixels (masked by hand) as a function of image row. The processed image is just the residual between the actual data and the best-fit quadratic. Automatic masking of non-road pixels should be possible either by interfacing with a road follower such as RALPH, or by performing road-edge detection in the reflectance image. We hope to show that laser reflectance provides adequate means for obstacle detec-

tion at longer ranges through more direct modeling of the sensor.

Figure 6. A laser reflectance image taken with the ERIM laser scanner. A block of wood 4x4 (approximately 4 inches high) was placed in the middle of the road at approximately a 20 meter distance. It is visible as the bright spot near the top middle of the image.

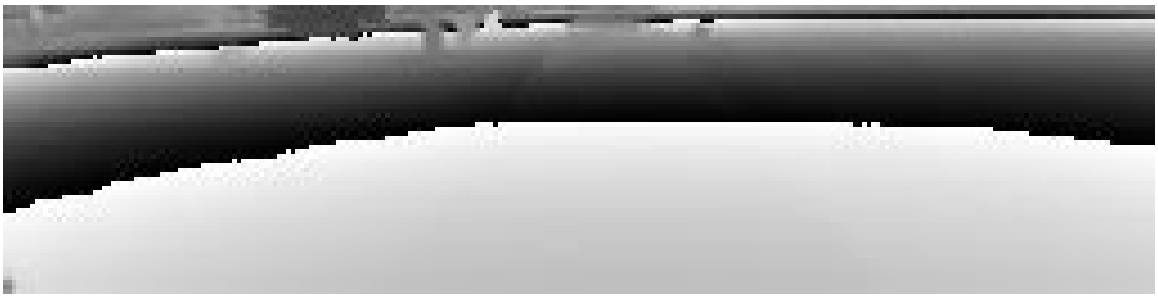

Figure 7. The ERIM range image of the same scene. The block of wood is not visible in this image. Darker pixels are closer and brighter pixels are farther with the exception of the change in brightness at the ambiguity interval of the sensor.

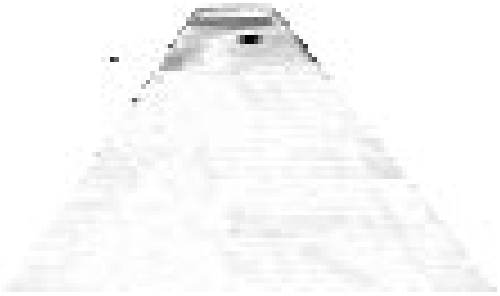

Figure 8. The processed reflectance image. Processing emphasizes those points in the image that don't fit the reflectance model well. Note that the non-road pixels were masked out by hand. The block of wood corresponds to the darkest spot.

Analysis of laser reflectance is complicated by its dependence on other factors such as range and material type of the sensed object. Assuming the diffuse component of the laser reflection is Lambertian, we can model the laser reflectance signal with the following relation:

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

where $\mathbf{P}_{return}$ is the power in the returning laser pulse, ρ is the actual surface reflectance ($0 \le \rho \le 1$), θ is the angle of incidence of the beam with the surface and **z** is the depth [23]. Assuming a vertical object with square corners, θ should be anywhere from zero to 45 degrees (due to rotations about the z-axis). For horizontal surfaces at distances over 20m, however, θ should be between 87 and 90 degrees and should provide a much smaller return.

The digital intensity value returned by the sensor may not be proportional to the returned power, however. The ERIM laser, for example, passes the returned signal through a log amplifier to reduce the dynamic range of the signal before sending it through the analog-to-digital converter [25]. Given the transfer functions of the filtering stages, however, we can calculate a reflectance value which is proportional to the power returned. Using either the range signal from the laser or the current road model, we can find **z**, allowing us to correct for range dependencies. Although we showed a flat world model is inadequate for detecting small obstacle by their height, it may prove adequate for estimating the range for normalization of reflectance. Correcting for the range by assuming a flat world is simple and has the advantage that it will also emphasize or brighten obstacles since points on above-ground obstacles are closer to the sensor than the road model would predict [10]. By also predicting the orientation, calibrating for constants, and having some knowledge of the reflectance (given a well-maintained road), we could a priori generate an expected reflectance image. Comparing actual data with expected data would reveal potential obstacles.

Detection of obstacles would then require the following steps (see Figure 9). First using the laser signal processing model, a laser power image could be calculated from the intensity image. Next, assuming a flat ground plane, we can normalize the power for distance and angle effects by using the following relationships:

$$\gamma = \alpha + ky$$

$$z = \frac{h}{\sin\gamma}$$

where α is the inclination of the sensor, γ is the inclination of an individual laser ray, **k** is a row factor, and **h** is the height of the sensor. We can produce a reflectance image from the power image by the following equation:

$$\rho(x, y) = \frac{z^2 P(x, y)}{\cos\theta} \quad \text{and for ground plane points,} \quad \theta = \frac{\pi}{2} - \gamma$$

Note that for vertical obstacles, **z** is smaller than predicted and cosθ is much larger than predicted, so pixels corresponding to vertical obstacles will be given a very high ρ value. Any points with unusually high ρ values should correspond to vertical obstacles. A diagram of

the processing is shown below.

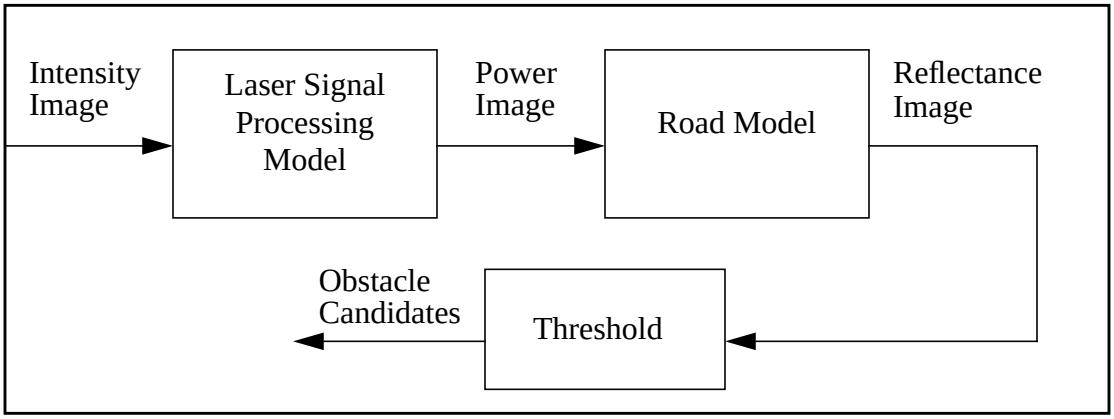

Figure 9. A laser intensity processing flow diagram. Given an intensity image, a laser signal processing model is used to obtain a value proportional to the returned laser power. A geometric road model is then used to normalize the power value for variations in range and angle of incidence to obtain a true reflectance image. Ground plane points should have acceptable reflectance values while obstacles should have much higher reflectance values.

There are several variations on this architecture depending on whether all obstacle candidates are passed to the vehicle controller or whether the system obtains road location information and filters out those obstacles which are out of the vehicle path. Road location information might be obtained from a road follower or found using laser intensity or reflectance data (before or after normalization for range). Although a flat world model is inaccurate when used in estimating road height as we showed previously, it is still reasonably accurate for purposes of distinguishing obstacles from road surface orientation. We expect that the difference in obstacle versus road intensity values will be large enough that a better road model is not needed. Some of the above steps may also prove to be unnecessary. In particular, it may not be necessary to normalize for angle of incidence variations over the road surface. Obstacles might be detected easily from a pseudoreflectance image (where pixel intensities are proportional to $\rho\cos\theta$).

The ease with which we can detect obstacles at long distances depends on numerous sensor parameters. Increased laser power improves the signal-to-noise ratio and increases the range (but decreases eye safety). Sensor angular resolution and beam size combine to determine overall sensor acuity. Pixel and frame rates effectively limit sensor resolution for mobile applications. To gain insight into how the factors affect system performance, we will test the method described above with up to 3 different lasers -- the ERIM, Riegl, and Z&F (see Table 2 on page 12 for sensor specifications). The ERIM and Z&F lasers are both continuous wave devices, and the Riegl is a pulsed laser. Although the Riegl has the longest range (over 100 m), the pulsed laser may prove unsuitable since amplitude measurements of a pulsed signal are likely to be less reliable than that of a continuous wave. We will design a distinct laser signal processing model for each of the lasers. Although some of our current laser sensors have a larger instantaneous field-of-view (IFOV) than required by our acuity constraint (see Table 2 on page 12) of 0.1 degrees (as listed in Table 1 on page 4) to detect a 20 centimeter obstacle at 65 meters, we believe that we can scale down our experiments appropriately with these sensors. Even without the necessary acuity, a laser with a large spot size may be able to detect an obstacle since a larger spot averages the reflectance signal over a larger area. Although the averaged signal provided by a laser with a spot size larger than the target provides us with a weaker signal difference, it will provide adequate coverage and may prove

sufficient.

If the above method does not prove sufficient, more sophisticated modeling of the road and/or laser may be necessary. Empirical evidence will then be collected to verify and refine the relationship between the return laser power, the range, the reflectance, and the angle of incidence.

Table 2: Laser Parameters

	ERIM	Riegl	Z&F
Average Laser Power	150 mW	1.6 mW	40 mW
Signal Type	Continuous	Pulsed	Continuous
Wavelength	820 nm	904 nm	780 nm
IFOV (Beam divergence)	0.5 deg	3.3 mrad ~ 0.2^o	0.3 mrad ~ 0.02^o
Resolution	256x64	variable	variable
Total FOV	80^o x 30^o	360^o x 30^o	360^o x 30^o
Laser Pixel Rate	92 kHz	12 kHz	500 kHz
Frame Rate	2 Hz	0.04 - 4.0 Hz	0.04 - 4.0 Hz
Intensity resolution	8-bit	8-bit	16-bit

6. Stereo Vision

There has been a large amount of research into using stereo vision for obstacle detection. Off-road and indoor work makes up the bulk of it, but more recently, there have been stereo vision systems designed for highway obstacle detection.

The first step in traditional stereo methods involves calibrating the camera setup. This involves finding the epipolar geometry of the two or more cameras. Once the epipolar geometry is known, the disparity at a pixel may be computed by searching along its corresponding epipolar line in the second image for the “best” match. Since in the general case, a rectangular window in the first image does not match to a rectangular region in the second image, the images may be rectified before matching so that the epipolar lines correspond to the image scanlines. This has the added benefit of making the search easier to implement. A match value is then found by computing the sum of squared differences (SSD) of the pixel values (or a similar metric) between a region around the point and the corresponding region around the hypothesized matching point. The point on the epipolar line that results in the lowest SSD value is chosen as the best match. In this way, the disparity may be estimated at each pixel in the image. The disparity map often needs to be filtered using left-right-line-of-sight (LRLOS) consistency checks, smoothing, and/or other techniques [22]. From the disparity map, obstacles might be detected in a number of ways -- by building elevation maps or looking for edges in the disparity map, etc. Knowing the precise position of image points in cartesian world coordinates requires some additional calibration.

There are a variety of pitfalls with this basic method. First of all, it is computationally

expensive -- it takes O(**N*M*K*L*D**) operations where **N** and **M** are the dimensions of the image, **K** and **L** are the dimensions of the window, and **D** is the number of disparity levels searched per pixel. Second, reliable matches can only be found in textured areas of the image. The algorithm may find wildly wrong disparity values in bland areas of the image (fortunately, these potential errors are generally detectable by looking at the shape of the SSD curve). Third, using the SSD metric over a window assumes that the pixels in the window are at a constant depth. Even after image rectification, if the pixels in a window lie at multiple depths, a rectangular window in the first image will not map well to a rectangle in the second image. This may result in poor estimates of the disparity value. It is the second and third concerns that lead to a conflict when deciding how large the matching window should be. The window must be large enough to contain some texture, but small enough so that it doesn't cover multiple depths. It is this conflict that led Kanade and Okutomi to propose the variable window size stereo algorithm[14]. The last drawback is that the standard algorithm treats each pixel individually in the final disparity estimate. It is poor form that the algorithm makes the implicit assumption that adjacent pixels are close to one another (through the use of the SSD window), yet does nothing to either verify or assure this in the final disparity map. This leads a number of artifacts that make it difficult to rely on individual pixel values. With regards to performing road-based obstacle detection, using this standard approach would be too slow and would ignore a lot of approximately known scene structure.

6.1. A Model-Based Approach for Stereo Obstacle Detection

We propose a model-based stereo vision approach for road-based obstacle detection. It is based on the fact that roads are nearly planar for the areas we will consider in a given image. The first step in the obstacle detection algorithm is to generate a disparity map. First we will consider the flat world assumption, and then we will assume an incremental polyhedral model. In both cases, we will show that, over the region of interest, disparity is a linear function of image row of the form $d = a + by$, where **y** is the image row and **a** and **b** are the offset and slope of the function.

6.1.1. Flat World Assumption

We start with the basic stereo equation:

$$z = \frac{bf}{x2 - x1} = \frac{bf}{d}$$

where **z** is the distance to the world point, **b** is the distance between the image centers or baseline, **f** is the focal distance, and **d** is the disparity (measured in units of distance).

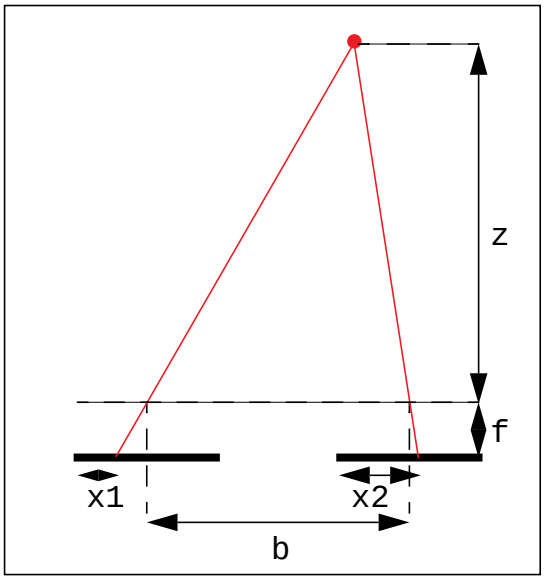

Figure 10. A point in the world, when imaged by two parallel cameras with focal lengths `f` and separated by a baseline `b`, appears in locations `x1` and `x2` in the images. The distance to the point, `z`, may be determined from the disparity `d` = `x2` - `x1`.

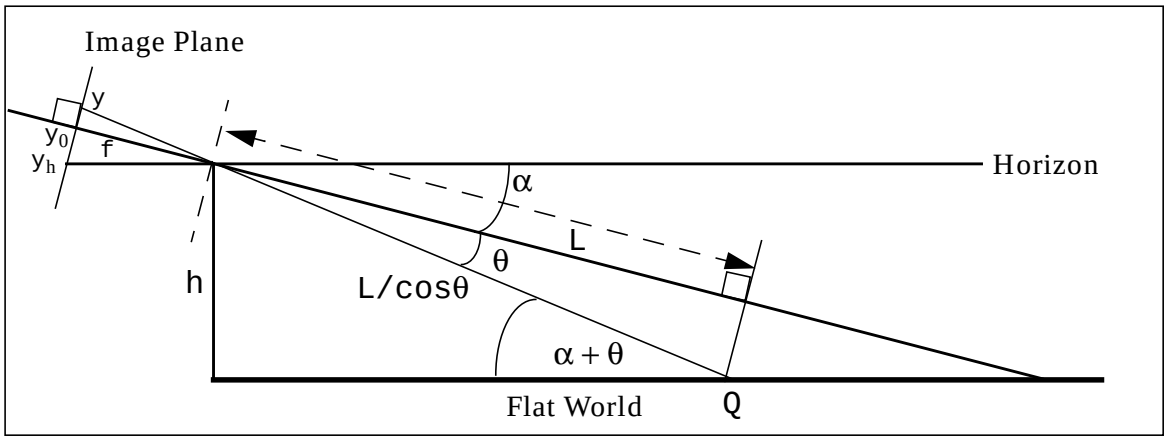

Figure 11. A camera with an angle of inclination α images a ground point `Q`. Point `Q` lies along a ray at an angle θ to the optical axis.The perpendicular distance from the focal point to `Q` is `L`.

Assuming a camera height of **h** above a flat world and a perpendicular distance **L** to a point **Q** on the ground plane as in Figure 11, we can write the following equations:

$$\sin(\alpha+\theta) = \frac{h}{L/(\cos\theta)} = \frac{h\cos\theta}{L} \qquad d = \frac{bf}{L} \qquad \tan\theta = \frac{y-y_0}{f}$$

Solving for disparity as a function of distance between the image row **y** and the image center,

we have for image rows below the horizon (disparity is zero for rows above the horizon):

$$d = \frac{bf\sin\alpha}{h} + \frac{b\cos\alpha}{h}(y - y_0)$$

$$\text{equivalently} \quad d = \frac{b\cos\alpha}{h}(y - y_h) = k(y - y_h)$$

where $\mathbf{y}_h$ is the row that images the horizon (the location of zero disparity). Although the units of disparity and image row must be measured in distance for the former equation, the latter equation holds for disparity and image rows measured in pixels. Although most of the variables in the above equation may be considered constant, vehicle motion will result in some pitching motion in the camera that will affect α, the angle of inclination. Bohrer, et. al., claims vehicle pitch variations of +/- 5 degrees are typical [6]. To find the right disparity model for a flat world given the constant parameters, we need only search for α, the angle of inclination. Alternatively, we note that for small α, cos α is nearly constant so that the slope, **k**, of this function will be nearly constant, and we may search just in the offset term ($\mathbf{y}_h$ in the latter equation). This allows us to calibrate for the slope and offset of the linear function via search using initial images without solving or calibrating for any of the individual parameters such as **b**, **f**, and **h** (although **b** and **h** are easy to measure approximately and may be used to compute an initial estimate of **k**).

The implicit assumptions here are that the images have been rectified so that the epipolar lines correspond to the scanlines (just as in the standard stereo algorithm) and that camera roll about the optical axis is negligible so that points on a scanline are at a constant depth.

6.1.2. Polyhedral Earth

Some roadway obstacle detection methods have made the flat earth assumption [6],[11]. As we showed earlier, this is a poor model since deviations from the flat earth may be as large or larger than the obstacles we wish to detect. Now we consider the case where the road has vertical curvature as shown in Figure 12. We have previously shown that a best-fit plane can adequately model the area of the road we are inspecting.

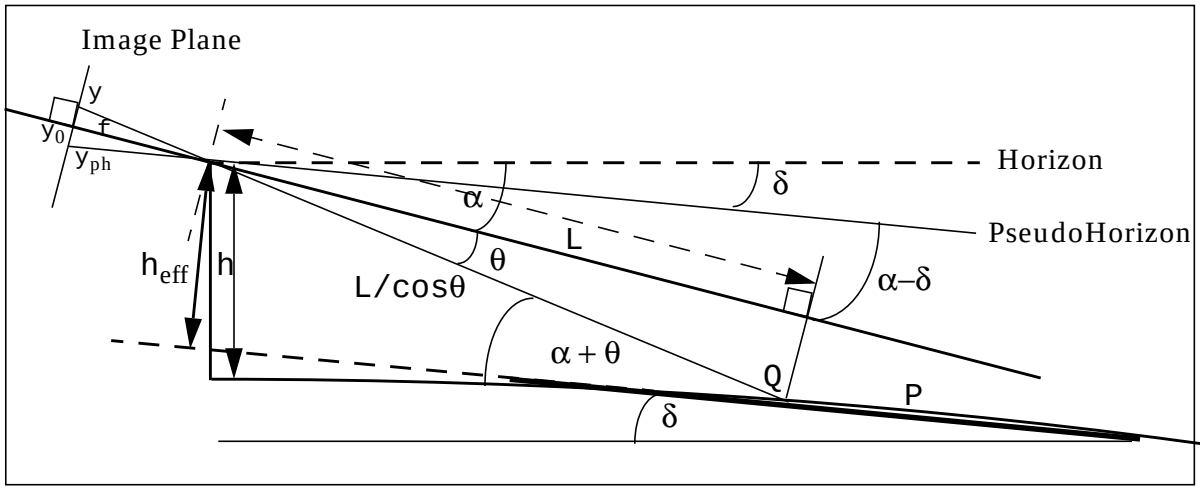

Figure 12. We have replaced the flat world from the previous figure with a curved road. We approximate the road by a plane P at an inclination δ for all rows of interest y in the image. The camera's

effective angle of inclination becomes $\alpha-\delta$ and we create a new variable h_{eff} for the effective height of the cameras.

If the best planar approximation **P** to the region of interest is inclined at an angle δ, our effective horizon is also tilted by this angle, since the two must be parallel so that they meet at infinity. We can see that the geometry is the same as in the flat world case with two slight modifications. First, our angle of inclination must be replaced by $(\alpha-\delta)$, the angle the optical axis makes with the planar road and the pseudohorizon. Second, the height **h** must be replaced by the effective height h_{eff}, which is the perpendicular height of the focal point above the plane **P**. Thus, we can rewrite our disparity equation:

$$d = \frac{bf\sin\alpha_{eff}}{h_{eff}} + \frac{b\cos\alpha_{eff}}{h_{eff}}(y - y_0) \quad \text{where} \quad \alpha_{eff} = \alpha - \delta$$

$$\text{equivalently,} \quad d = \frac{b\cos\alpha_{eff}}{h_{eff}}(y - y_{ph}) = k'(y - y_{ph})$$

where y_{ph} is the row that images the pseudohorizon (the location of zero disparity). To find the best fit disparity model, we now have to find two parameters, α_{eff} and h_{eff}, both of which may vary over time. Since h_{eff} is a function of δ and other constants, the parameter search could be done in α and δ instead. However, it is probably simplest to search (at each time step) for the composite slope and offset parameters (k' and y_{ph} in the latter equation) to avoid unnecessary calibration of the constants. Since δ should vary slowly, we would expect the slope of the disparity function, k', to vary slowly while the offset, y_{ph}, should vary more rapidly (due to camera pitch variation). Both of these terms may be bounded by further geometric calculation and empirical results.

The simple linear relation of disparity to image row allows us to warp the right image so that points on the road match the left image. If we assume a polyhedral earth, we need to search during operation (at each time step) for both the proper slope and offset parameters for the disparity function. This is done by calculating the normalized correlation (or other metric such as SSD) over all pixels in the left image that have a match in the right image (as determined by the disparity function) for a range of k' and y_{ph} values (see Figure 13). The disparity function $d = k'(y - y_{ph})$ that produces the maximum correlation between the entire images is chosen as the correct one.

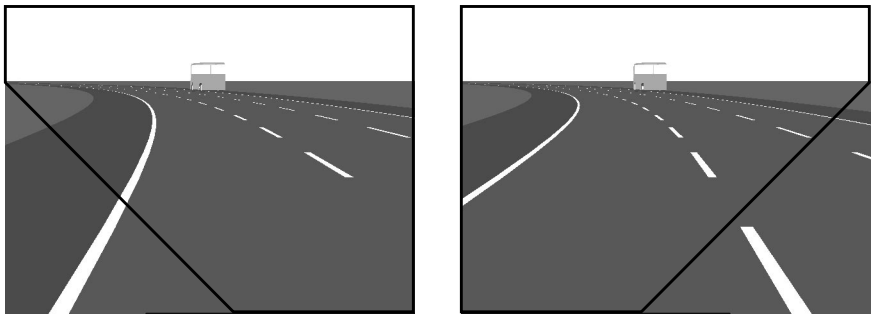

Figure 13. This shows a simulated left-right stereo pair. For all pixels on the road, the pixel locations in the two images may be related by the equation $x2-x1 = k'(y-y_{ph})$. The first step of our algorithm, after image rectification, performs a normalized correlation between the two images for a variety of k' and y_{ph} values. The regions outlined in black contain the pixels that are included in the correlation for a given k' and y_{ph} (in this case, the correct values).

This method has a number of advantages over the traditional stereo method for calculating the disparity map. First, it has a reduced computational complexity of O(**M*N*Y**) where **M** and **N** are the image dimensions and **Y** is the number of values searched in the horizon for the flat world case, or O(**M*N*D**) where **D** is the number of disparity levels that has to be matched for each image row (which is dependent on the variability in the slope and offset parameters) for the polyhedral earth case. More importantly, it can use the features of the entire image to produce a match at every pixel while still maintaining the proper relationship of depth between adjacent pixels (for pixels which actually lie on the road). It uses the known structure of the road to produce a consistent disparity map of the form $d = k'(y - y_{ph})$.

This work is most closely linked to that of Koller, Luong, and Malik which uses essentially the same road model [16]. They formulate the problem a bit differently, however, allowing only one degree of freedom in the effective angle of sensor inclination, α_{eff}. Their model does not allow for changes in $\mathbf{h}_{\text{eff}}$ which are caused by road curvature [16]. Their approach also differs in how they distinguish obstacles from road, our next topic. They detect obstacles using the model-based disparity map, while we use a comparison of expected versus actual image intensities.

6.1.3. Detecting Obstacles in the Difference Image

Given the disparity map from the previous step, we can subtract a warped version of the right image from the left image. Ideally, points that lie on the ground plane should "disappear" (have a value of zero) in this difference image, making the obstacles easily detectable. In practice of course, this does not happen. Problems can occur for a number of reasons. First, discretization of the images will keep us from warping one image to exactly match the other (see Figure 14), and since we warp each image once to rectify them and warp the right image to match the left, we may introduce a number of re-sampling artifacts. Second, intensities of correctly matched points are generally not the same since camera gains may be different and the images are noisy. Finally, our model parameters may not be exact due to discretization of the parameter search space, noise, or large obstacles which violate our assumptions. Given these problems, we need to develop a method for determining whether a non-zero area in the difference image is caused by an obstacle or other artifacts.

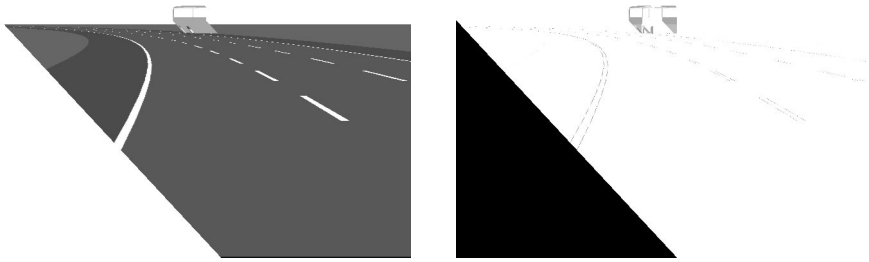

Figure 14. If we warp the right image by the proper disparity function, we obtain the left image of this pair which looks very similar to the left image of the original stereo pair (the white area at the bottom-left indicates pixels where no match was possible). If we subtract this image from the original left image we obtain a difference image (we have inverted the image for clarity). Notice that most of the ground plane disappears except at some edge boundaries. The vehicle ahead of us shows up easily.

One approach to determine whether a non-zero area in the difference image is caused

by an obstacle is to attempt to calculate the actual disparity for the non-zero region and compare that with the expected disparity. However, this has the problem that individual disparity calculations for a finite window may be unreliable and occasionally, drastically wrong. Instead, our method models the intensity values themselves to determine whether a non-zero area is caused by an obstacle. We warp the right image by the disparity function (see Figure 14). Linear interpolation is performed between adjacent pixel values for non-integral disparities. For every pixel in the left image, we compare a rectangular window around it to a rectangular window located at the same place in the warped image. The size of the window will be a trade-off between noise immunity, sensitivity, and computational expense. Large matching errors indicate potential obstacles.

To define "large matching error", we use an approach taken from Matthies [20] and model the signals from the stereo images as offsets of the same intensity signal with noise added to each image:

$$I_l(x) = I(x) + n_l(x)$$

$$I_r(x) = I(x + d(x)) + n_r(x)$$

where $\mathbf{I(x)}$ is the intensity signal and $\mathbf{d}$ is the true offset or disparity between images $\mathbf{I}_l$ and $\mathbf{I}_r$, and $\mathbf{n}_l$ and $\mathbf{n}_r$ are noise (in pixel brightness). The noise is assumed to be Gaussian with zero mean and equal variance, σ^2, in each image. A single row in the warped version of the right image may be described by:

$$I_w(x) = I(x + \Delta x) + n_r(x)$$

where $\Delta\mathbf{x}$ is how far we're off from the correct match (this should be small). Given a signal with a single disparity (a single image row with no obstacles), and using the sum of squared differences (SSD) as our error metric (which we can obtain by squaring the pixel values in the difference image), we have:

$$SSD = \sum_{x_i} (I_l(x_0 + x_i) - I_w(x_0 + x_i))^2$$

$$\Rightarrow SSD = \sum_{x_i} (I(x_0 + x_i) + n_l(x_0 + x_i) - I(x_0 + x_i + \Delta x) + n_r(x_0 + x_i))^2$$

where $\mathbf{x}_0$ is the point around which we are matching, and the $\mathbf{x}_i$ are the offsets for the pixels in the window. For small $\Delta\mathbf{x}$:

$$I(x + \Delta x) = I(x) + \dot{I}(x)\,\Delta x$$

where $\dot{I}$ is the intensity gradient, measurable from the image. By combining the noise terms, simplifying and using the linearization, we obtain:

$$SSD = \sum_{x_i} \left(\dot{I}(x_0 + x_i)\,\Delta x + n(x_0 + x_i) \right)^2$$

With the best disparity model for an image, we should be close to the correct disparity value for ground-plane pixels. If errors are only caused by sampling problems, then $\Delta x \leq 0.5$. Given a noise model and gradient information for the original images, we can use this error model to calculate whether the error for a region around a given pixel is too large to be caused by sampling and noise differences or whether it is caused by an obstacle. We can use

a similar metric based on the sum of absolute differences (SAD).

So far, we have ignored the effect of different camera gains in the model. One way to reduce the effect of different gains is to apply a Laplacian of Gaussian filter to the images. The Gaussian removes some of the high frequency information some of which is due to noise, and the Laplacian removes the low frequency data such as that caused by different gains. Unfortunately, this also greatly reduces the information content in the image. Perhaps a better error metric that would not require pre-filtering is normalized correlation (as we already used to find our optimal disparity function). Choosing an appropriate error metric for both finding the best disparity model and for finding obstacles in the difference image will be one of the important elements of the thesis work.

6.1.4. Stereo Architecture Summary

Before running the system, calibration of the epipolar geometry of the stereo cameras must be performed to calculate the image rectification matrices. There are a number of methods proposed for this [12],[26]. We use the 2-plane method to calculate the epipolar geometry [24] and choose rectification matrices which minimize image distortion [26]. When running the system, the images are digitized and then rectified. After rectification, the system searches for the disparity function that best matches the right image to the left. The parameters of this function may be used to constrain the search space on the next image pair. Once the optimal disparity function is found, the right image is warped to match the left and a window about each pixel in the left image is matched to the same window in the warped version of the right image. An error metric (still to be determined) is then used to determine whether a poor match corresponds to an obstacle or not.

Potential difficulties that might hinder detection include: sensitivity to unmodeled parameters such as camera roll about the optical axis, high camera noise and camera gain variability, and our road model being led astray by the presence of large obstacles.

6.2. Experimental Setup

We have built a strapdown stereo jig for use on the CMU Navlab 5 research vehicle. The cameras are attached to a stiff, rigid bar which should keep epipolar geometry nearly constant. Currently, the jig has a 1.25 m baseline, and is located at a height of approximately 1.0 m. The focal length of the lenses is 16 mm. This configuration was chosen to balance a number of trade-offs. A large baseline was chosen to maximize the depth resolution. Normally, baselines are kept smaller to allow for more overlap between the camera fields of view. However, at the long distances we are considering, the baseline has little effect (compared to the field of view) on the overlap. The depth difference (and hence detectability) between an object and the road increases as the sensor height is reduced -- at the limit, where the sensor height equals the object height, the camera is looking at the object or infinity. Thus, the jig was placed on the hood of the vehicle rather than the roof. This height should be enough to see over vertical crest curves according to the highway design specifications previously mentioned. Finally, to obtain an ample field of view, we chose the minimum focal length that would provide detectability according to the equation:

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

where **f** is the focal length required to obtain an **n** pixel difference in disparity between the

road and an obstacle of height **p** at a distance **L** given a camera baseline **b**, height **h**, and physical cellsize **c**. For a 1/2" 640 x 480 CCD, a 10 cm object at 65 meters will result in over a 3-pixel change in expected disparity for 16 mm lenses. Our lenses provide a 20^{o} horizontal field of view which allows us to see around highway curves without steering the camera jig. This field of view also will enable us to cope with significant variations in camera pitch due to vehicle motion. If this configuration does not provide adequate detection, we may double the lens focal length at the cost of halving our field of view.

7. Conclusion

Highway obstacle detection is a difficult problem and crucial to the future of such programs as AHS. Despite extensive work on the problem, it remains unsolved. By comparison with the cross-country problem, we see that a road model is needed. While a flat world model may be sufficient for detecting large objects, we have shown it is not accurate enough for small obstacle detection. We must not choose poor models for the sake of convenience. Only by examining sensor design and actual road design can we rigorously verify our models' accuracy.

We have presented two methods for highway obstacle detection. Preliminary results with ERIM laser reflectance look promising, and better modeling of the sensor internal processing should improve results. Two additional lasers, the Riegl and the Z&F, will be tested and modeled. A practical model-based stereo method has also been presented. Based on our calculations presented here, we expect that a locally planar road model will be sufficient to detect 20 cm obstacles at 60 m. However, further experimentation and calculations will be used to verify the model's accuracy and sensitivity to other parameters such as sensor roll.

8. Expected Contributions

We expect to develop a system that will be able to detect small static obstacles (20 cm high) at high speeds and long distances (60 m). We hope to demonstrate the feasibility of using laser reflectance as a sensor modality for obstacle detection. A number of elements will be generated to reach this goal:

- analysis of how road geometry affects obstacle detection (by examining highway design manuals)
- laser reflectance characterization and modeling (through analysis of processing circuitry)
- analysis of error metrics for stereo vision applications (through signal modelling and empirical results)
- a design for a practical, model-based stereo method for roadway obstacle detection
- a novel method for using laser reflectance for obstacle detection

9. Schedule

Spring 1997	Thesis proposal presentation Characterize Riegl and Z&F lasers Evaluate error metrics for model-based stereo method Test detection algorithms on ERIM data
Summer 1997	Evaluate stereo method on roadways Refine road models for stereo Continue evaluation of Riegl and Z&F lasers Design laser processing models for Riegl and Z&F lasers
Fall 1997	Integrate detection methods with a method for finding road boundaries such as RALPH (a road-following system). Begin testing of integrated detection systems.
Spring 1998	Finish testing of integrated detection systems. Perform algorithm sensitivity analyses. Begin writing thesis.
Summer 1998	Finish writing thesis. Defend.

10. 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]Asada, M. Building a 3-D World Model for a Mobile Robot from Sensory Data. *Univ. of Maryland Technical Report*, CAR-TR-332, CS-TR-1936, October 1987.

[4]Bishop, J. R. Intelligent Travel: The Automated Highway System. *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS-3)*, 1993.

[5]S. Bohrer, 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]Hartley, R. and R. Gupta. Computing Matched-epipolar Projections. *Proceedings of the Conference on Computer Vision and Pattern Recognition (CVPR '93)*, 1993.

[13]Heisele, B. and W. Ritter. Obstacle Detection Based on Color Blob Flow. *Proceedings of the Intelligent Vehicles '95 Symposium*, 1995.

[14]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.

[15]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.

[16]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.

[17]Krotkov, E., R. Hoffman. Terrain Mapping for a Walking Planetary Rover. *IEEE Transactions on Robotics and Automation*, Vol. 10, No. 6, Dec. 1994.

[18]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.

[19]Langer, D. J. Rosenblatt, M. Hebert. A Behavior-Based System for Off-Road Navigation. *IEEE Transactions on Robotics and Automation*, Vol. 10, No. 6, Dec. 1994.

[20]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.

[21]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.

[22]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.

[23]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.

[24]Oda, K. Calibration Method for Multi-Camera Stereo Head for NavLab II. Internal CMU Document, 1996.

[25]Pont, W. Basic Characteristics of the 3-D Sensor. Erim Staff Report, Range and Reflectance Processing Workshop, Warren, MI, December 1987.

[26]Robert, L., M. Buffa, and M. Hebert. Weakly-Calibrated Stereo Perception for Rover Navigation. *Proceedings of the International Conference on Computer Vision (ICCV)*, 1995.

[27]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.

[28]Ross, B. A Practical Stereo Vision System. *Proceedings of Computer Vision and Pattern Recognition (CVPR)*, 1993.

[29]Ruichek, Y. and J.-G. Postaire. Real-Time Neural Vision for Obstacle Detection Using Linear Cameras. *Proceedings of the Intelligent Vehicles '95 Symposium*, 1995.

[30]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.

[31]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.