CTTR 67-95/01

Great Lakes Center for Truck and Transit Research

Characterization of Road-Edge Markings in Support of Road-Departure Prevention Systems

Gregory E. Johnson
University of Michigan
Transportation Research Institute
Ann Arbor, MI 48109

Karl C. Kluge
University of Michigan
Ann Arbor, MI 48109

Robert D. Ervin
Transportation Research Institute
University of Michigan
Ann Arbor, MI 48109

October, 1995

2201 UMTRI Building, 2901 Baxter Road, Ann Arbor, MI 48109-2150 313/936-1064
Associated Universities: Central State University • Michigan State University
Michigan Technological University • University of Michigan • Wayne State University

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
GLCTTR 67/95-01		88435

4. Title and Subtitle	5. Report Date
Characterization of Road Edge Markings in Support of Road-Departure Prevention Systems	October, 1995
	6. Performing Organization Code
7. Author(s)	8. Performing Organization Report No.
G E. Johnson, K. C. Kluge, R. D. Ervin	UMTRI 95-42
9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
University of Michigan Transportation Research Institiute 2901 Baxter Road Ann Arbor, MI 48109-2150	11. Contract or Grant No. DTRS92-G-0005 ~~DTRSG-92-0005~~
	13. Type of Report and Period Covered
12. Sponsoring Agency Name and Address Great Lakes Center for Truck and Transit Research 201 UMTRI Building 2901 Baxter Road Ann Arbor, Michigan 48109 ITS Research Center of Excellence University of Michigan Ann Arbor, MI 48109	**Final Report**
	14. Sponsoring Agency Code

15. Supplementary Notes

Supported by a grant from the U.S. Department of Transportation, University Transportation Centers Program

16. Abstract

The visual properties of painted lane markings were characterized through a field measurement exercise using a specialized trailer apparatus. The class of markings examined consisted of the continuous white stripe defining the boundary between the right lane and the shoulder of the road. Stripe width, reflectance, and contrast with lane and shoulder regions were quantified for over 6000 video images taken along a 600 mile sample of Michigan highways. Data indicate that slightly less than 82 percent of the stripes surveyed had widths in the range of 8 to 15 cm., with wider stripes outside this range being more common than narrower stripes. Stripe reflectance exhibited a bimodal distribution with primary modes around 28 percent and 60 percent reflectance values. Application of the data set to an existing road-follower algorithm is discussed.

17. Key Words	18. Distribution Statement
	None

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
None	None	32	None

Contents

1.0 INTRODUCTION

Many electronic systems to prevent road-departure, facilitate the lane-keeping task, or otherwise support a control functionality based on lateral placement require that road-edges be continuously sensed by some type of imaging technology [1,2-8,9-18,21-25]. Even though many systems in this class rely on charge-coupled device (CCD) imaging technology and the reflective properties of the contrasting lane marker, little work has been done to inventory real roads so as to meaningfully characterize visible road-edge markings as they actually exist. This project was a field measurement exercise by which image-based data were collected and processed to describe the qualities of retroreflective road-edge markings as found on Michigan roads. These characterizations of stripe reflectances, contrast, and physical dimensions are based on a modest sample of interstate highways within the state.

Current system designs for road departure warning and intervention functions have relied primarily on CCD cameras and ambient illumination to sense the retroreflective paint stripes. Subsequent image processing using a variety of processors with a wide range of capabilities yields an estimate of the roadway geometry ahead. The quality of the estimate is a combination of the sensor quality, illumination, stripe quality, road surface covering, and degree of sophistication of the processing algorithm. The stripe quality is one aspect that has a direct bearing on detectability and is the primary subject of this investigation. Other factors, such as illumination and environmental elements were held fixed or characterized where possible.

The project's three objectives reveal the pragmatic focus on the characteristics of actual in-service highway lane markings. These were:

- to assemble a trailer-mounted instrumentation system with which to collect data on those properties of road-edge markings that are pertinent to the performance of roadway imaging systems based on the visible spectrum;
- to collect such data on a statistically meaningful sample of Michigan highways;
- to process the data and produce distributions in a manner that gives guidance for the development of road-imaging systems.

The imaging and processing effort yielded several distributions of properties describing the sampled stripes. The information can be used to test limits of current algorithms or designs, or to provide a basis for investigating the state of technology required to be consistent with the current state of striping.

The rest of the document is as follows. Section 2 presents the instrumented platform and discusses the various aspects of its design. Section 3 covers data collection including a discussion of the scope of data included in the study. In section 4, we present the data processing steps, including data calibration, used to produce the results presented in section 5. Conclusions are stated in section 6.

2.0 INSTRUMENTED PLATFORM

2.1 Mechanical description

The trailer was designed around a utility trailer kit. The axle assembly and surrounding frame structure were slightly modified to allow for an enclosed observation cell and a rear platform for generators. Photo 2.1.1 shows the trailer mounted behind a tow vehicle.

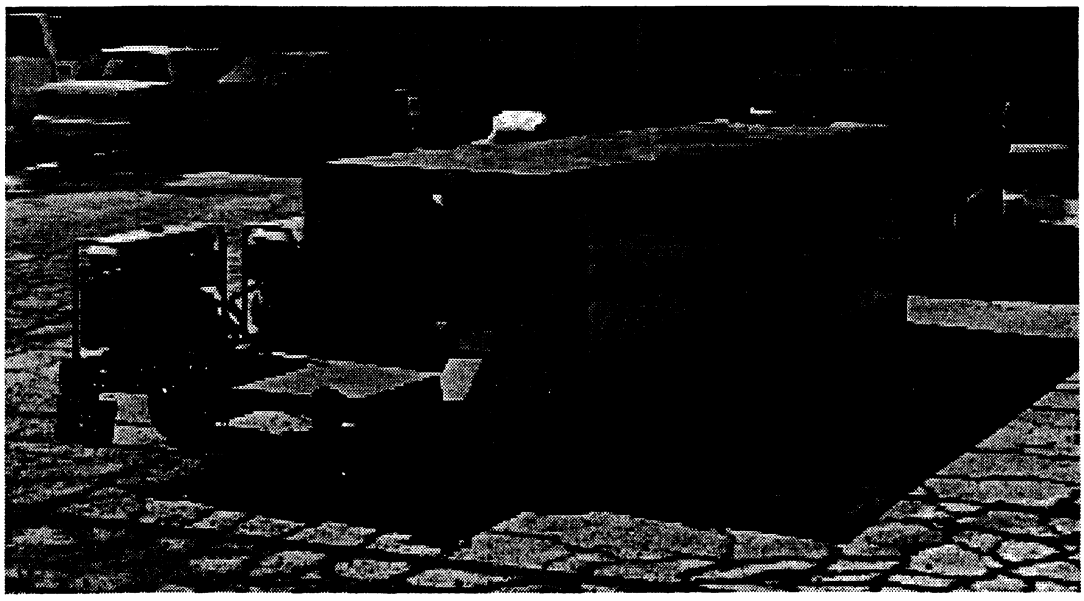

Photo 2.1.1. Lane Marker Characterization Trailer with tow vehicle

The length of the trailer was determined by the need to obtain a relatively long illumination space and a shallow angle of view. This was intended to provide diffuse ambient illumination, and allow to the imaging of the stripe at roughly the same viewing angle that would be used by a low- or mid-mounted camera on a passenger vehicle or truck. The use of an enclosure also eliminates shadows on the portion of the road being imaged, and provides control over the illumination. The scene captured by the test platform is shown in photo 2.1.2 below.

Shown in Photo 2.1.1, the large box section of the trailer contains the observation cell - a cavity that is open only on the bottom for viewing the road markings under conditions of controlled illumination. The trailer is comprised of tubular steel and a riveted aluminum skin that forms a shroud to block ambient illumination. The entire underside of the shrouded observation cell has no lateral cross members and thus provides unobscured sensing of the roadway below, as seen in Photo 2.1.2. The interior of the trailer is painted white to aid reflection of the light sources, which are mounted at the top of the trailer. Stiff brushes and a rubber flap extend downward from each trailer side to make contact with the roadway. This provides better screening of the observation cavity from exterior illumination. Two generators are mounted on the rear of the platform, providing AC power for the onboard systems.

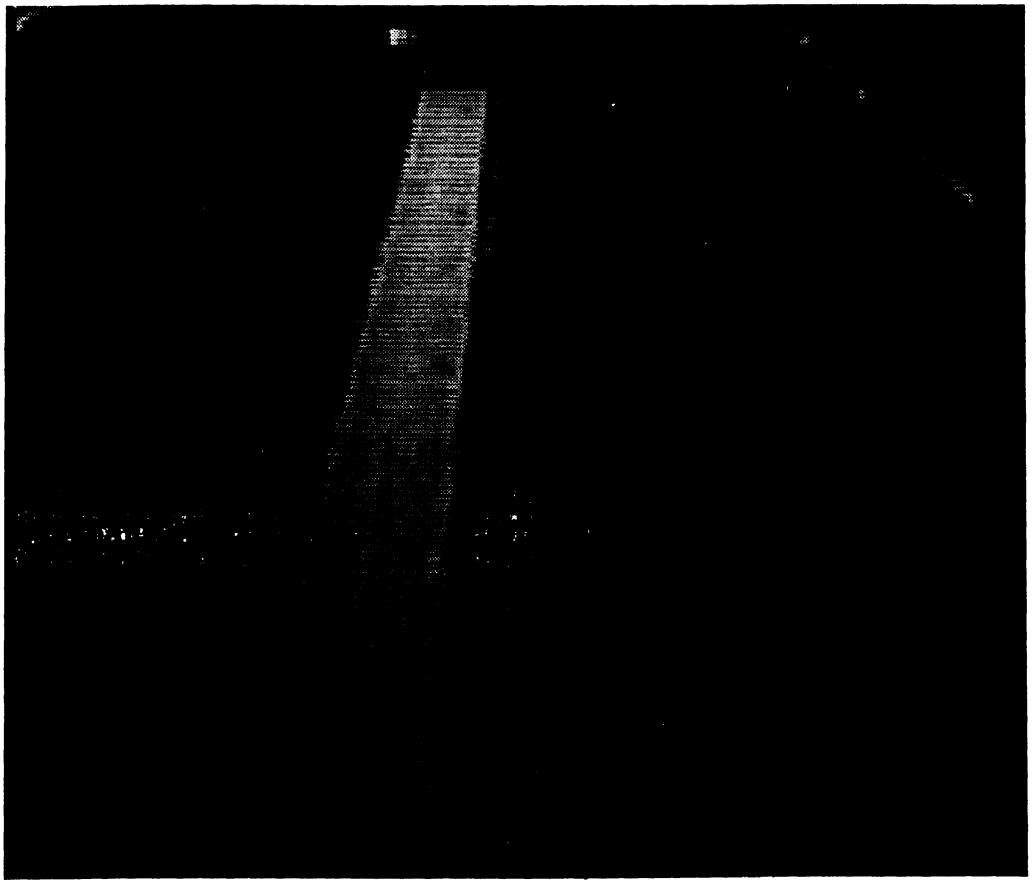

Photo 2.1.2. Image as seen with the trailer - example data

The imaging system installed consisted of a Phillips monochrome CCD module with a 12.5mm Computar f1.8 manual focus/iris lens. The camera had electronically adjustable gain, black level (offset), and shutter speeds. The speed of the shutter was fixed to 1/250 seconds, and the gain/offset output adjusted for maximum dynamic range of the

calibration card with respect to the lighting system installed. These parameters, as well as the iris and focus, were held fixed for the duration of data collection.

2.2 Illumination

A broad spectrum diffuse source was desired for good approximation of a constant ambient illumination. Although fluorescent lamps provide efficient, fixed spectrum illumination, they exhibit fast reaction times such that their intensity can vary greatly with the power delivered. This output fluctuation (present in everyday use, but unnoticed at 60Hz by the human eye) creates problems when using modest shutter rates with CCD cameras. An incandescent light source also receives alternating current for power, but has a much larger thermal mass and a longer reaction time. Thus, changes in current delivered to an incandescent lamp are effectively low-pass filtered by the filament's reaction time and a much more continuous illumination level is provided.

Reaction time became an issue in our choice of illumination sources because we chose a fast electronic shutter speed. Recording images at the industry video standard of 29.97 frames/sec while illuminating the scene with lights powered from a 60Hz AC supply, the varying phase relationship causes each frame to be snapped on a different portion of the AC sine wave, resulting in large changes in illumination if fluorescent lamps are used. The use of incandescent lights was an efficient solution to the illumination problem.

3.0 DATA COLLECTION

In this study the scope of data was limited to a representative sample of Michigan highways. Analog image data were collected using the instrumented platform described above. The data from the camera were recorded using a high quality VCR on 1/2" VHS format video tapes. A 'valid-data' signal was also recorded on one of the audio tracks in the form of a sine-wave, as well as the SMPTE (Society of Motion Picture and Television Engineers) time code standard.

The 'valid-data' signal consisted of a 200mV, 650Hz sine wave, which was recorded onto the audio track of the video tape. The frequency was the desired vehicle speed multiplied by 10, in mph (i.e. 650Hz = 65mph, 550Hz = 55mph). The signal was 'on' or being recorded when the trailer was positioned over a region of roadway such that the

paint stripe should be visible and detectable within the observation cell. Cases in which the 'valid' signal was off (nothing recorded on the sound track, but video still being recorded) include portions of the roadway environment not covered by the scope of the study, such as entrance and exit ramps, construction zones, and temporarily marked regions. Also, regions in which the driver was not able to keep the trailer over the road edge stripe due to road debris or obstacles were marked as invalid data.

Data collection progressed on a 15-minute cycle. Each consisted of a seven minute session of driving with the trailer on the stripe, followed by eight minutes of driving normally, without collecting data. This fairly lax duty cycle was selected in order to minimize the fatigue on a test driver. Several short trials around the Ann Arbor region were conducted early in the study to test equipment and procedures. Sample data were also processed to exercise the processing software.

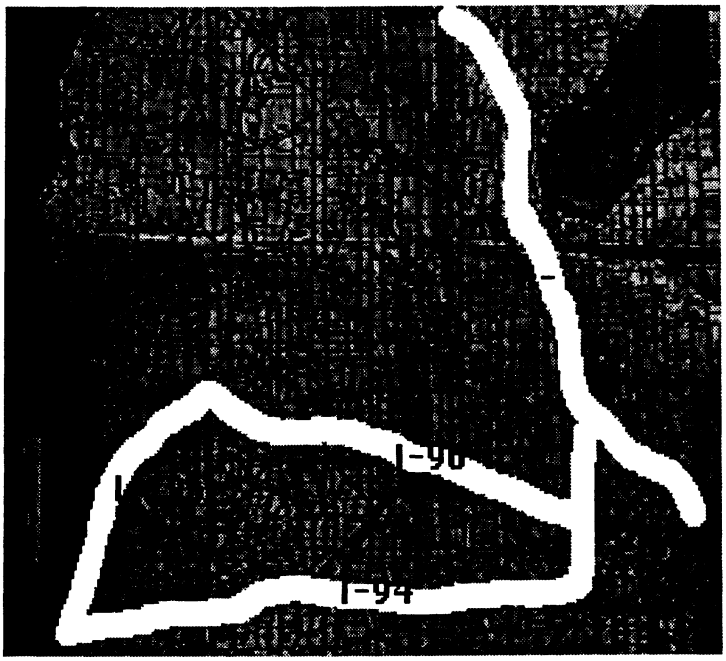

Figure 3.1. Michigan map indicating highways surveyed

For collection of the data presented in this report, two distinct routes covering more than 600 miles were driven in the summer of 1995. The route is indicated in figure 3.1 and includes a small stretch of US-23N, I-96W, I-196S, I-94E, and the lower portion of I-75, both North and South bound sides.

4.0 DATA PROCESSING

Data were digitized from the analog video tapes using the built-in frame grabber on a Macintosh Quadra 840AV. The frame grabber's parameters, such as bit depth and image quality, were held constant for all images. The digitization involved an automated sampling procedure in which the computer controlled data tape playback, sample selection (via monitoring the 'valid-data' signal on audio inputs), and provided a nominal 2Hz sampling interval. Over 6000 digitized images and the subsequent SMPTE address for each image were saved to disk. Note that the analog video tape is still the original medium and remains valid for redigitizing at any interval or digitizer settings. Since intensity calibration is also stored on the video tracks, it is possible to compensate for small degrees of tape degradation or drift in the video deck's analog components.

4.1 Detection Algorithm

Given the large volume of images to be processed, it was necessary to develop an automated technique to locate the edges of the painted stripe in the images digitized from the video tapes. Four features of the task drove the design of the computer vision algorithm to do the detection:

1) The algorithm needed to be robust. It needed to be able to find stripes that were very worn or otherwise degraded.

2) The algorithm should not use a strong model of the stripe appearance. Since part of the project goal was to characterize all actual stripes encountered — the good, the bad, and the ugly — the algorithm needed to be able to detect a broad variety of faint, overpainted, wide, and narrow stripes.

3) The algorithm did not need to run in real time. Since the data analysis is done off-line in the lab, processing times of several seconds per image were acceptable.

4) The use of the enclosure and the limited area imaged created a situation where the image intensity distribution should be essentially bimodal, with peaks corresponding to the average road and stripe pixel intensities.

Given these constraints, an algorithm was chosen based on the use of intensity histograms to select a threshold value that divides the image pixels into "stripe" and "background." The stages of the image processing are illustrated in Figure 4.1.1.

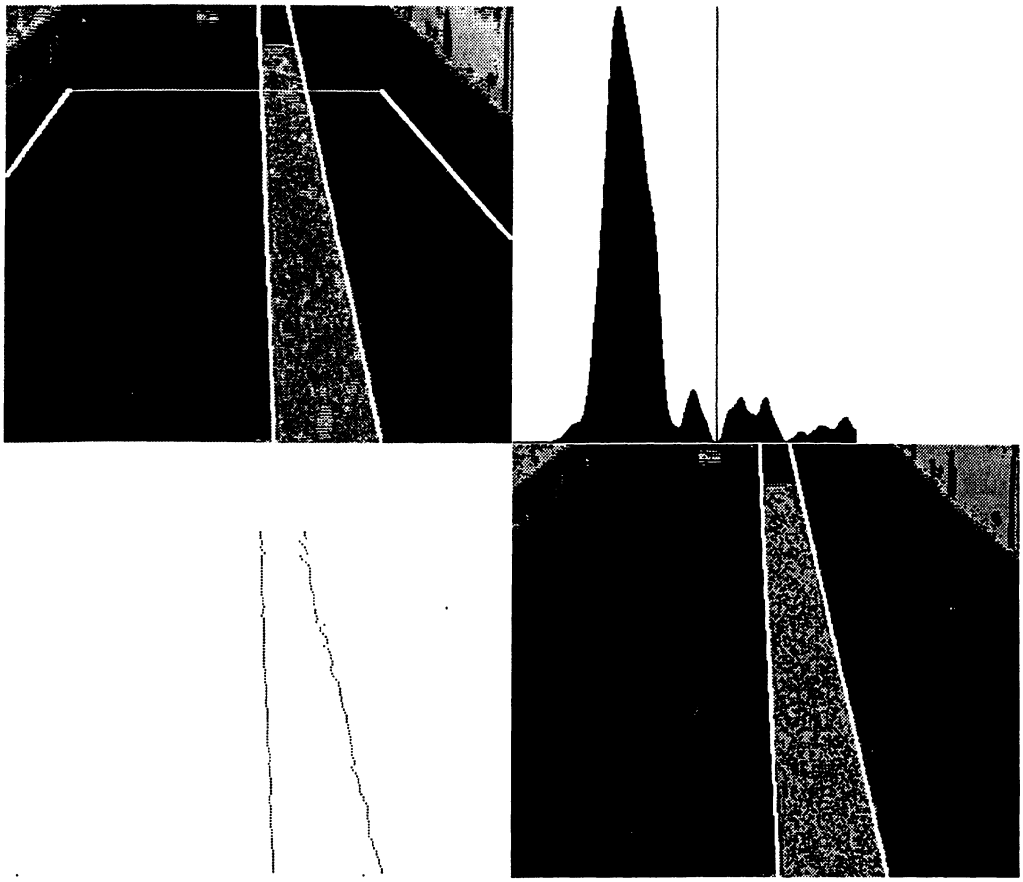

Figure 4.1.1. stages of image processing

The upper left quadrant shows the raw image with boundary and segmentation lines drawn on it. Since the location of the enclosure walls is constant from image to image, the image processing ignores the region that corresponds to the front and sides of the enclosure. The boundaries of the area ignored at the top and sides are drawn in white.

The rest of the image is used to compute a histogram — a curve that shows the number of pixels (vertical axis) that have a given image intensity value (horizontal axis). The histogram for this image is shown in the upper right quadrant of the figure. This histogram is used to identify the "best" intensity threshold to use to separate the pixels into "stripe" and "background." The heuristics used were inspired by techniques such as [20].

1) Local maxima ("peaks") and minima ("valleys") are located in the histogram.

2) Each possible triple combination of a peak, a valley to the right of the peak, and another peak to the right of the valley is examined.

3) A triple is rejected if

   a) placing the threshold at the valley would result in the stripe containing more than 50 percent or less than 5 percent of the image area; or

   b) the peak to the right of the valley is larger than the peak to the left of valley; or

   c) there is a higher peak between the valley and the left peak; or

   d) there is a higher peak between the valley and the right peak.

4) Among triples that aren't eliminated in step 3, the one with the largest difference between the height of the right peak and the valley is selected. The intensity threshold is placed at the level corresponding to the valley of the best peak-valley-peak triple. In Figure 4.1.1, the threshold selected is indicated by the vertical line drawn in the histogram in the upper right quadrant.

Once the image has been threshold, lines are fit to the edges of the stripe. The first step of this part of the algorithm locates the left- and right-most above-threshold pixels in each row of the image. These are shown in the lower left quadrant of the figure. Straight lines are fit to the left and right edge points subject to the constraint that the lines converge on the row in the image plane corresponding to the horizon (in this case, approximately 106 rows above the top of the camera field of view). The lower right quadrant of figure 4.1.1 shows threshold image, with the estimated stripe edges drawn in white. Once the stripe edges have been located, the stripe and road statistics for the image can be computed. The Intelligent Vehicles '95 Symposium paper appended to the report gives details of the evaluation of the accuracy of the program in locating stripe edges in an initial 615 image data set.

4.2 Intensity Calibration

Since the illumination was not precisely controlled (due to power and phase differences, broken light bulbs, dirt accumulation on lens, etc.) the ability to calibrate each image individually was desired. A gray scale calibration card was mounted at the front of the trailer shroud to provide a fixed-intensity reference for calibrating the images, visible at the top center of photo 2.1.2. The cal-card is shown below in photo 4.2.1.

Photo 4.2.1. enlarged view of intensity calibration card in trailer

The card contains 20 squares, each of a known reflectance. Considering that the CCD camera's input intensity is raised to a power (γ) by the output circuitry, a relatively simple relationship can be drawn [19]:

$$\%refl = a1*(I^{\gamma}) + a0 \qquad \text{(eqn 4.2.1)}$$

where 'I' is the image intensity gray scale.

It is known that the luminance of an object depends on the object's surface properties, as well as on the source intensity and location. The relationship is shown below

$$I = (\%refl) * \cos(\text{angle to light source}) \qquad \text{(eqn 4.2.2)}$$

Thus the design of the trailer illumination, while approximating ambient, presented a serious anomaly in that the source location(s) cannot be discounted since their location and interaction affect the intensity distribution. In fact, plots of intensities in figure 4.2.1 along longitudinal axes of the image indicate a parabolic curvature of intensity, whereas the intended condition was that of a diffuse, uniformly-distributed illumination.

While the simple relationship in eqn. 4.2.1 permits a relative calibration to the grayscale intensity card, the true road and marker albedo values cannot be derived. Breakage of individual bulbs and varying intensity prohibits a static correction as indicated via eqn 4.2.2, and dynamic correction was beyond the scope of this work. A more robust and diffuse illumination design may alleviate some of these drawbacks.

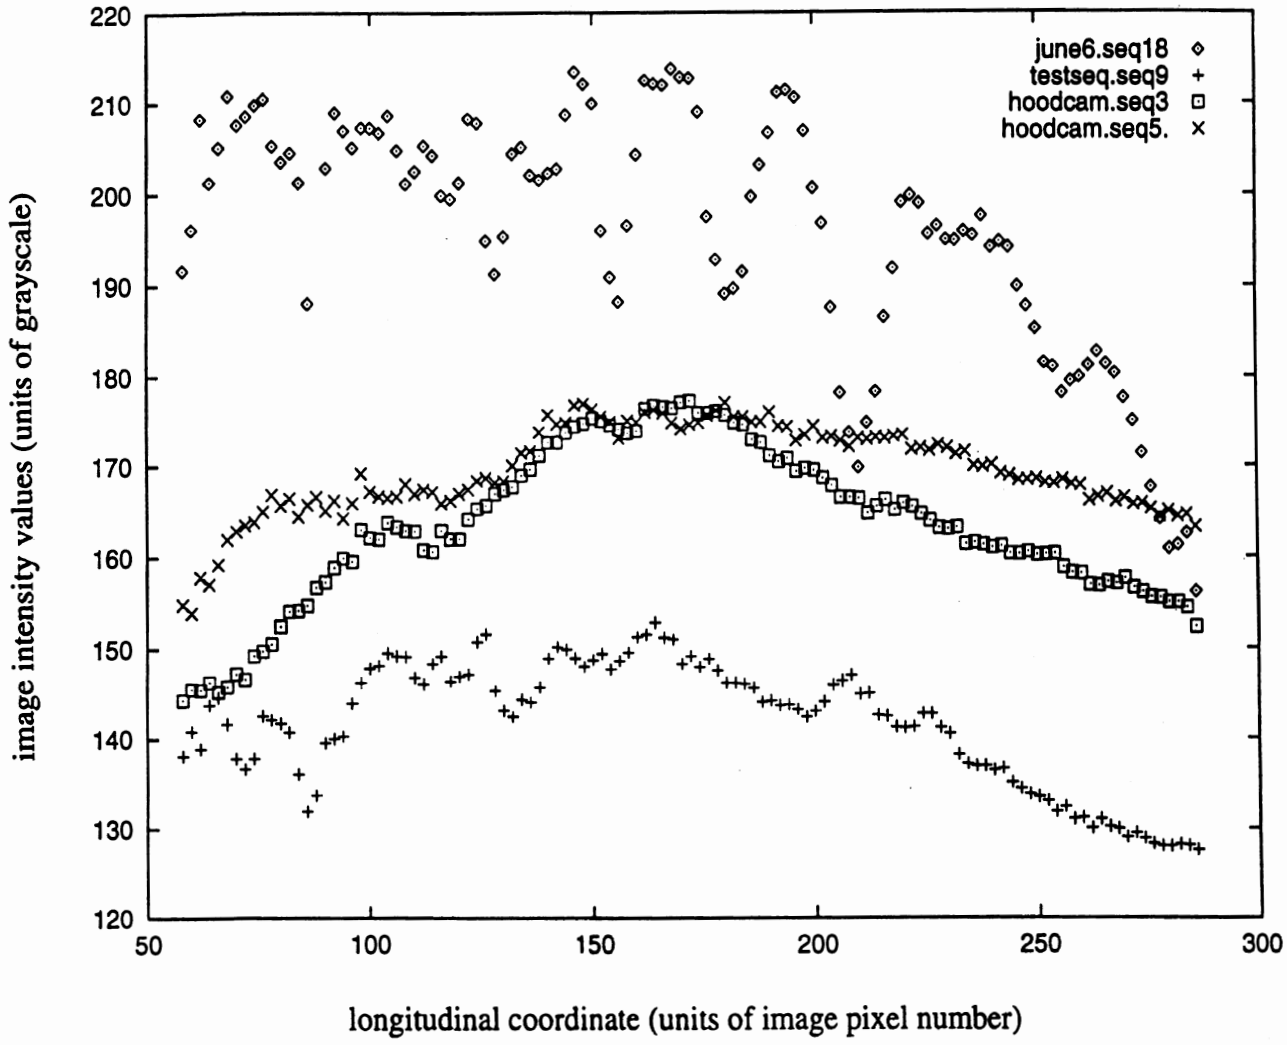

Figure 4.2.1. plots of intensity vs longitudinal coordinate

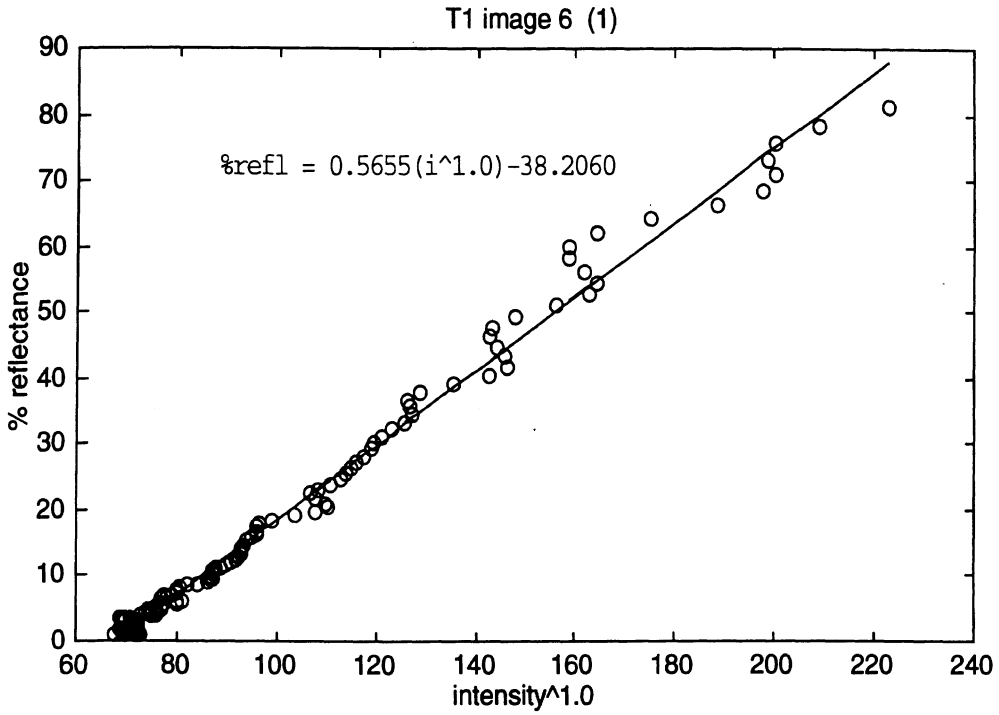

Figure 4.2.2. regression example of cal-card reflectances vs albedo values

Coefficients a0 and a1 in eqn. 4.2.1 were derived for individual images using the calibration card present in each image. The brightest square was automatically located, and the card's intensities between this square and that observed at a fixed distance to the right were taken as the cal reference. This was corrected for gamma and saturated values. Then a least-squares fit between the gamma-corrected gray scale values and the known card reflectances provided the calibration coefficients. An example fit is shown in figure 4.2.2.

4.3 Geometric Calibration

Geometric calibration of pixel locations to the roadway coordinates follows typical image coordinate conversion techniques [2,13,14]. First the camera/lens combination is calibrated to yield a factor relating width of pixels in the image to the angle subtended by the pixel. Pixels in the sensor are rectangular, and vertical and horizontal calibrations for the angles are different (note that figure 4.3.1 details only the lateral orientation).

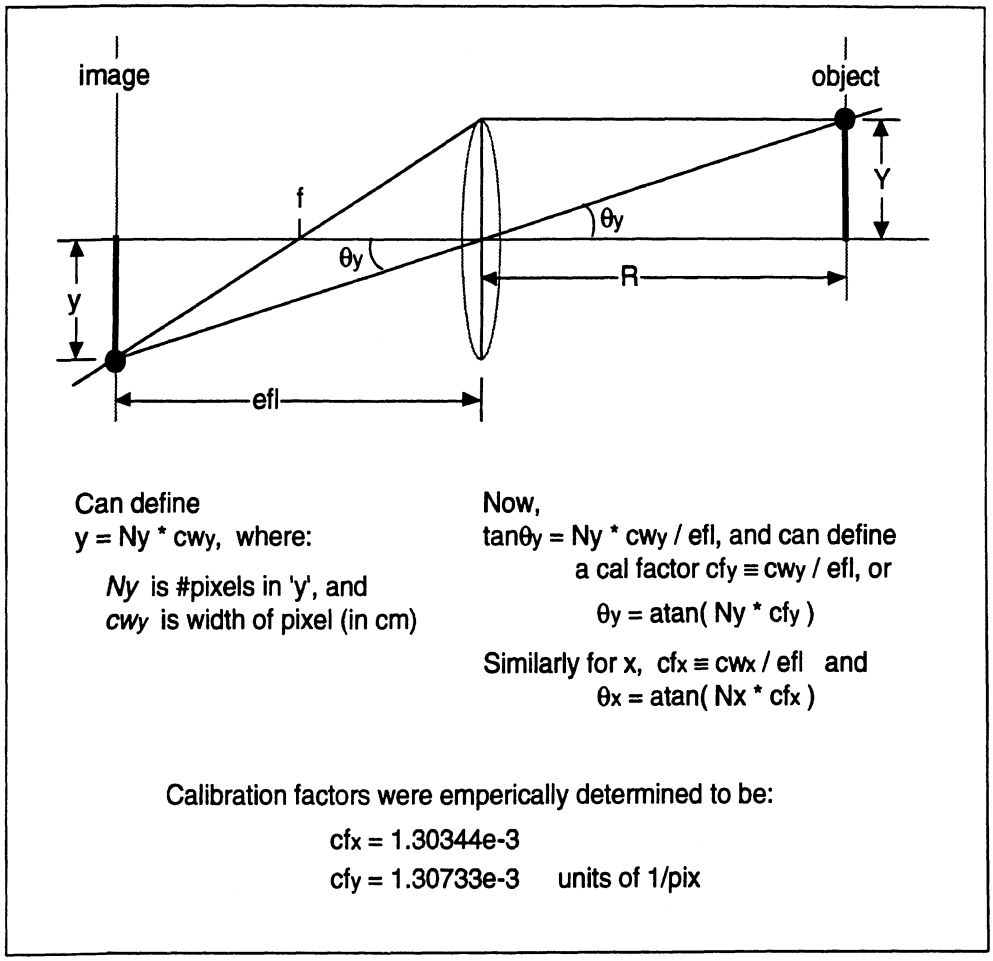

Figure 4.3.1. camera / lens calibration

Now the orientation of the sensor can be considered. It is assumed that there is a static depression angle (pitch) but no roll in the sensor. Since the camera is mounted on a suspended platform and used over rough surfaces the orientation will change to some extent. Nevertheless the scope of the study did not include direct measurement of pitch angles – believed to vary less than one degree. The coordinate system in figure 4.3.2 is Cartesian in the road plane, using standard sign notation with 0,0 located directly beneath the camera.

The geometric calibration is a relationship between the pixel/lens geometry and the camera location and orientation. Using the calibration factors, any pixel location in the image (Nx, Ny) can be transformed to a camera-centered coordinate system projected onto the roadway (Xc,Yc).

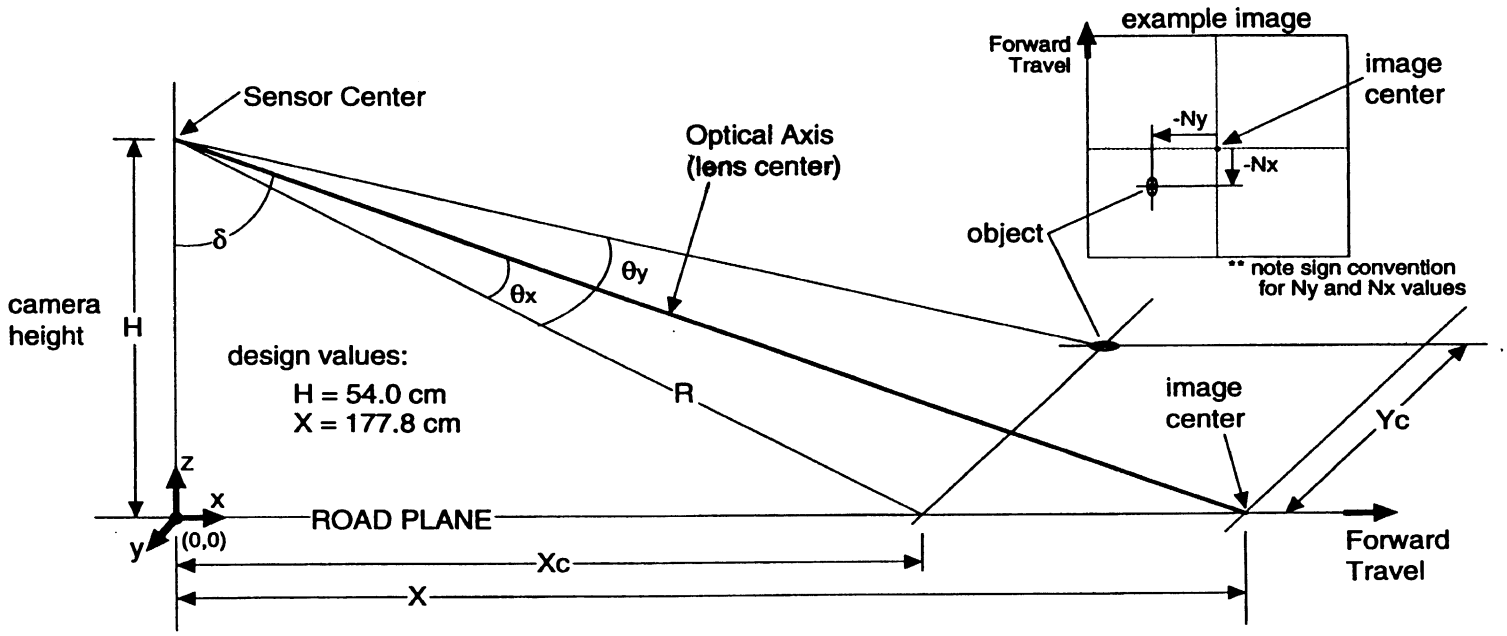

X - Calibration (longitudinal)

1. $\tan\delta = X/H \rightarrow \delta = \text{atan}(X/H)$
2. $\tan(\delta-\theta x) = Xc/H$
3. $\theta x = \text{atan}(Nx*cfx)$

combining 1, 2, & 3

$Xc = H\tan(\text{atan}(X/H) - \theta x)$

Y - Calibration (lateral)

1. $\tan\theta y = Yc/R \rightarrow Yc = R*\tan\theta y$
2. $R = \text{sqrt}((Xc^2 + H^2) / (1 + cfy^2 * Nx^2))$
3. $\theta y = \text{atan}(Ny*cfy)$

combining 1, 2, & 3

$Yc = \text{sqrt}((Xc^2 + H^2) / (1 + cfy^2 * Nx^2)) * (Ny*cfy)$

Figure 4.3.2. geometric calibration

5.0 RESULTS & DISCUSSION

Over 6000 images were digitized from the video tapes recorded during the trips described in Section 3.0. A video image consists of two *fields*. The even numbered rows of the image make up one of the fields, the odd numbered rows make up the other. The fields are recorded 1/60th of a second apart. At 60 mph, this corresponds to a distance of almost 1.5 feet traveled between the two fields that make up an image. As a result, each field is processed separately, resulting in over 12,000 data points.

Once the stripe edges were located in an image using the techniques described in section 4.1, the following statistics were computed:

- stripe width in cm.
- mean stripe reflectance
- mean lane reflectance
- mean shoulder reflectance
- mean road (lane + shoulder) reflectance
- standard deviation of stripe reflectance
- standard deviation of lane reflectance
- standard deviation of shoulder reflectance
- standard deviation of road reflectance

One of the local feature detection algorithms used to detect white stripes in the YARF road follower [13] illustrates the potential utility of this data. The processing stages of the feature detection are illustrated in Figure 5.0.1. The detection algorithm looks for the white stripe by averaging the image intensity along the expected direction of the stripe, and running a simple one-dimensional edge detector over the averaged signal. The stripe detector looks for a pair of edges with opposite contrast separated by the expected stripe width.

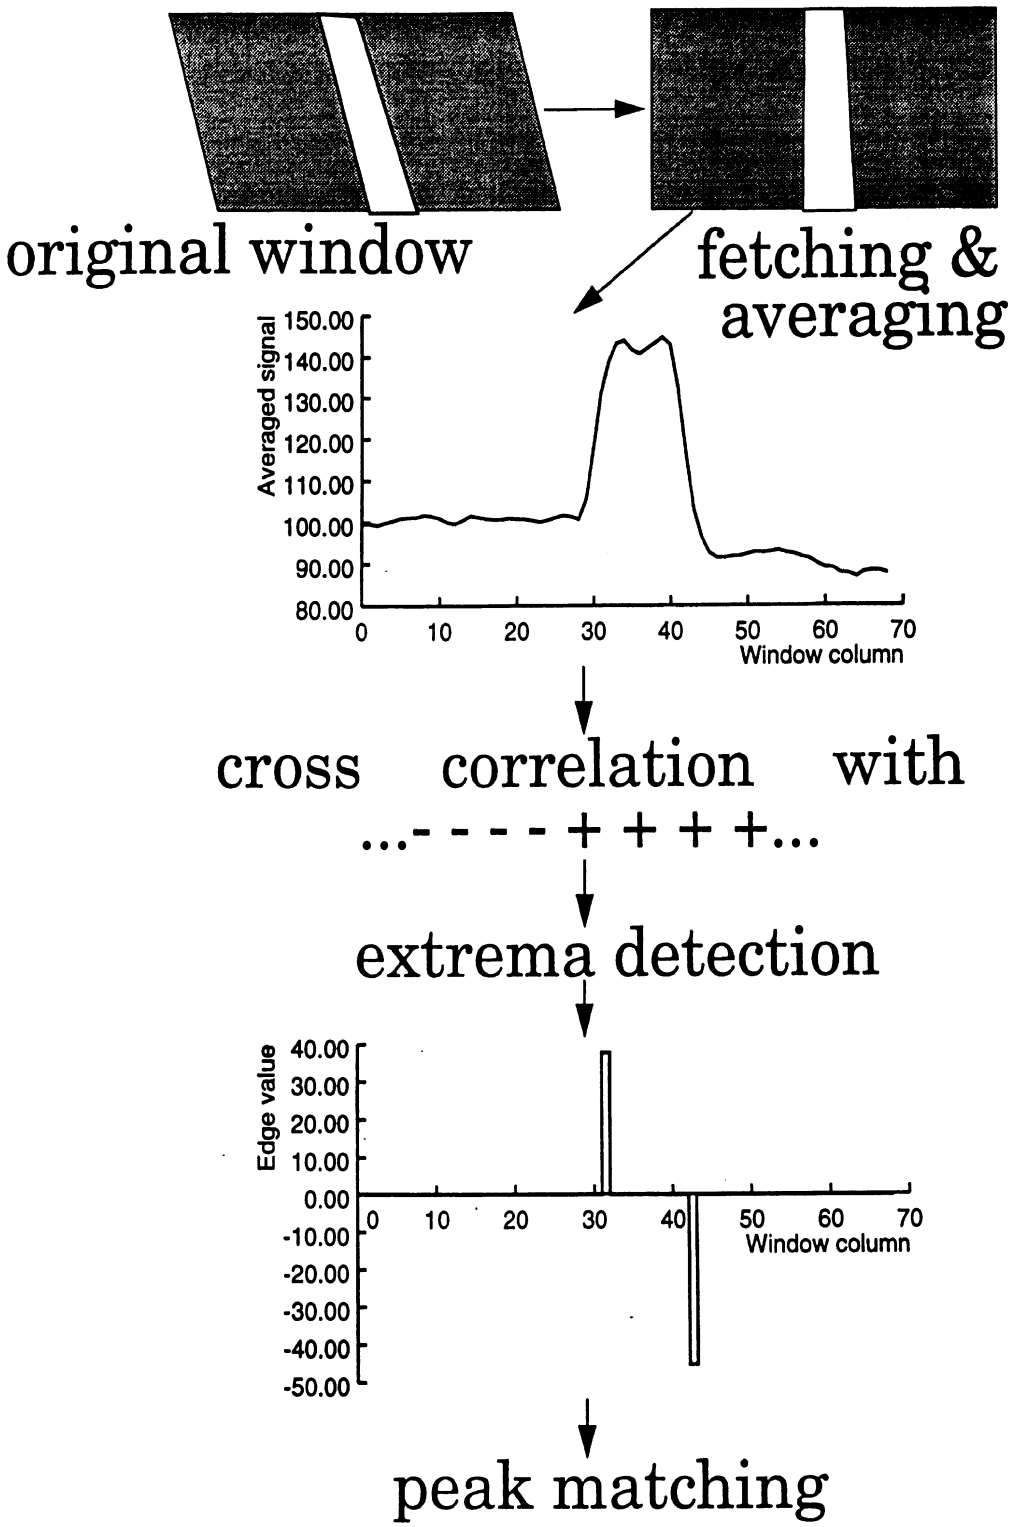

Figure 5.0.1. diagram of YARF's detection scheme, illustrating the similarity concept

The location of the stripe edges is detected by cross correlation of 1-D signal (the averaged image region in figure 5.0.1) with a simple 1-D edge filter. The filter uses a mask with a left negative half and right positive half. The width of the mask is equal to the expected width of the stripe. The weights used are +1 and -1 to allow fast cross correlation. The cross correlation value is normalized by dividing the raw value by half the width of the mask, producing a value that represents the intensity contrast at a given pixel.

The algorithm then locates extrema in the output of the cross correlation with the edge detector filter. This results in a list of candidate edges in the signal. This set of edges is searched for a pair of edges that meets the following three criteria:

- the edges have opposite contrast, with the signal brighter between the two edges,
- the edges are separated by a distance within a specified percentage of the expected feature width, and
- the absolute magnitudes of the two edges are similar (specifically, let Mp be the magnitude of the positive edge and let Mn be the absolute value of the magnitude of the negative edge — the pair is accepted if alpha < | Mp - Mn | / (Mp + Mn)). Mn will be proportional to the difference in stripe and shoulder reflectances, and Mp will be proportional to the difference in stripe and lane reflectances.

The tracker indicates that the feature was absent if no pair of edges in the window meets these criteria. The acceptable width tolerance around the expected value and the alpha threshold for the edge contrast similarity measure need to be specified. The data collected in this study permit those values to be selected in a way which reflects the actual distribution of those values on a significant set of Michigan roads.

5.1 PDFs and Cumulative PDFs

Figure 5.1.1 shows the distribution of stripe widths derived from the collected image data. The mode value is near 10 cm, or approximately 4 inches. The long right tail of the width distribution in consistent with the overspray of stripe, which often occurs during stripe repainting. Permitting the width to range from 8 cm to 15 cm (inclusive) would capture 81.9 percent of the stripes encountered in the sample for this study.

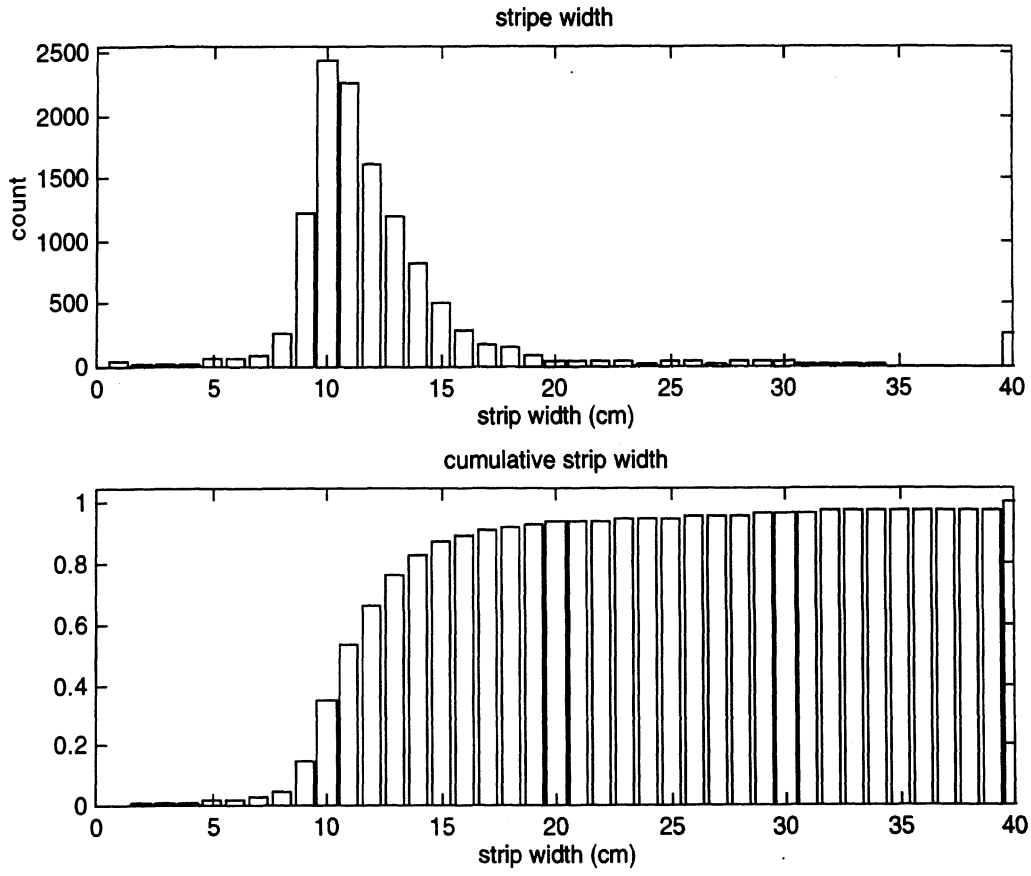

Figure 5.1.1. stripe width in centimeters

Figure 5.1.2 shows the distribution of the contrast similarity metric used. A value of zero indicates that the stripe has the same contrast with both the lane and the shoulder, with increasing values indicating decreasing similarity of contrast. Setting a threshold of alpha = 0.25 would capture 96.5 percent of the stripes encountered in the sample for this study. Thus, the false negative rate (number of stripes incorrectly not detected) can be controlled by setting thresholds based on the distributions derived from the data collected in this study.

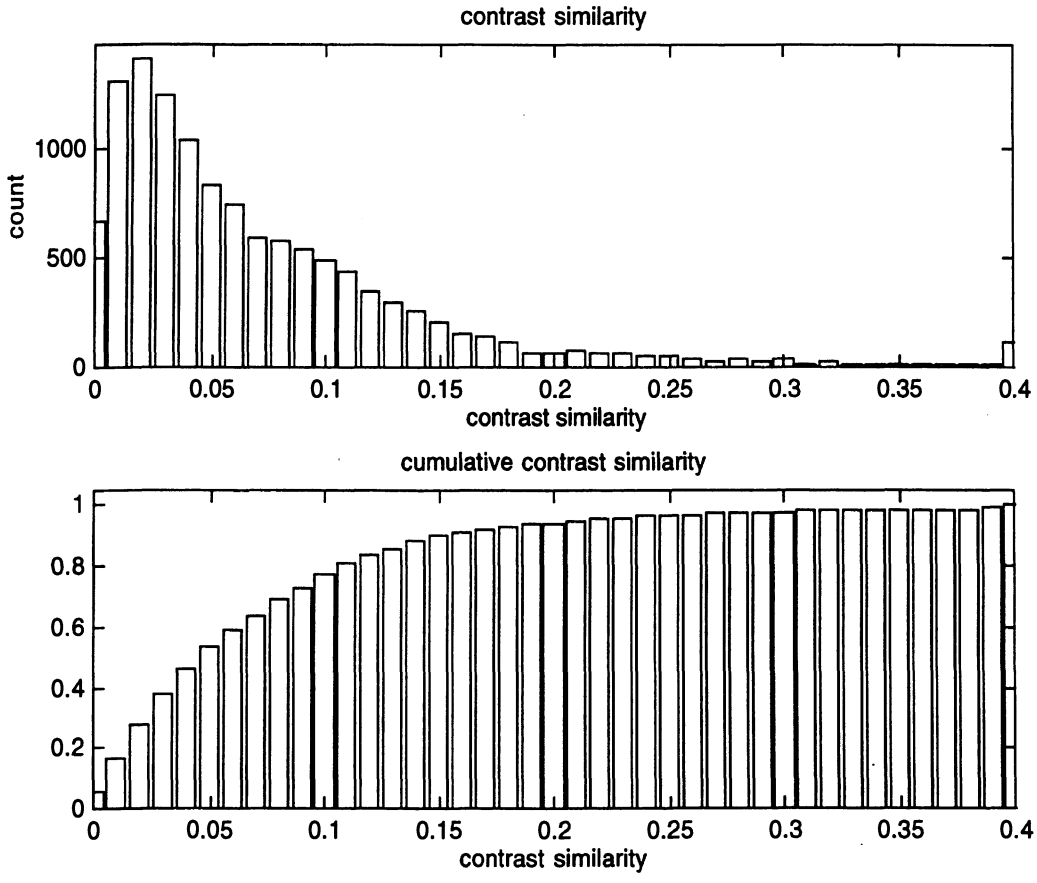

Figure 5.1.2. contrast similarity measure

A distribution of reflectance values of stripes in figure 5.1.3 indicate the broad bimodal distribution of intensities present. The differences in lane and shoulder intensities are also shown. Note both positive and negative values, indicating a spectrum of regions where the road is brighter than the shoulder and where the opposite is true. Note that the highest bin in the difference measure is non-zero. This is consistent with the similarity metric in which the largest mode was also non-zero.

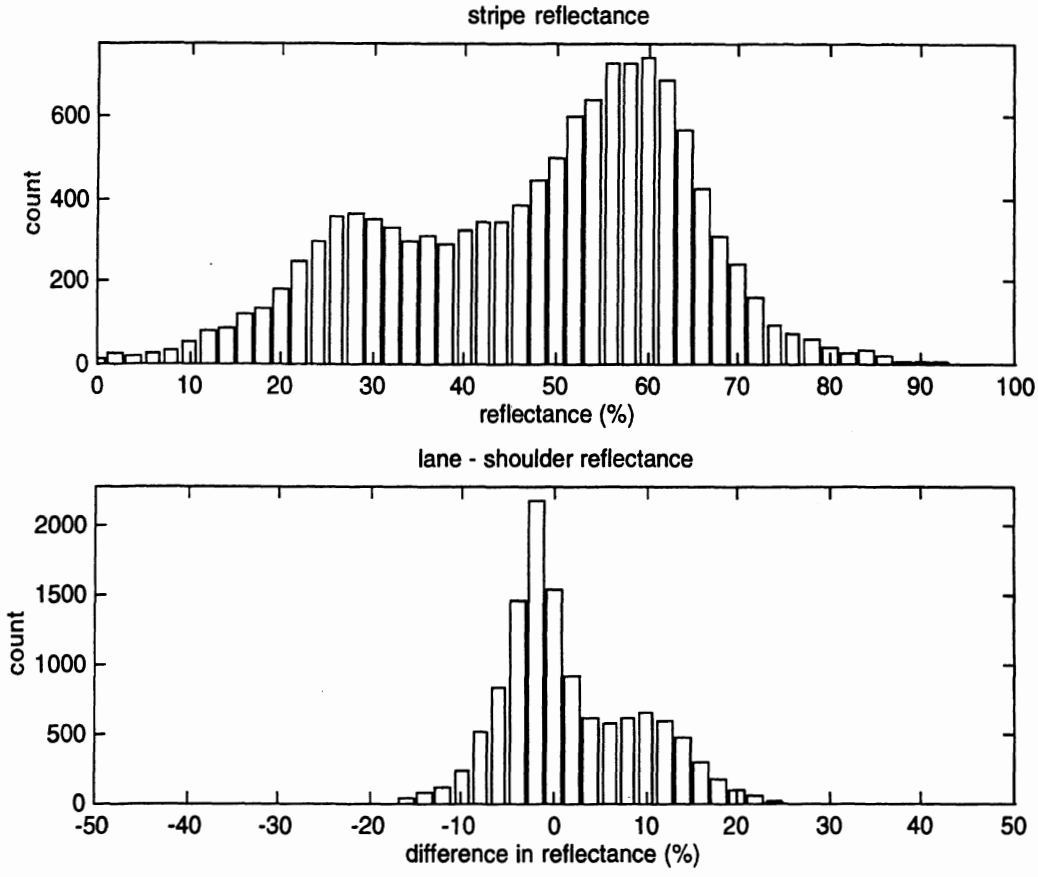

Figure 5.1.3. reflectance distributions

5.2 Two-Dimensional Distributions

Density plots of two variables provide insight into the decision space mapped by the two factors. Thresholds may have certain interactions, or dependencies upon other variables. Following is a density plot which suggests how similarity and stripe width (two of YARF's threshold parameters) may interact.

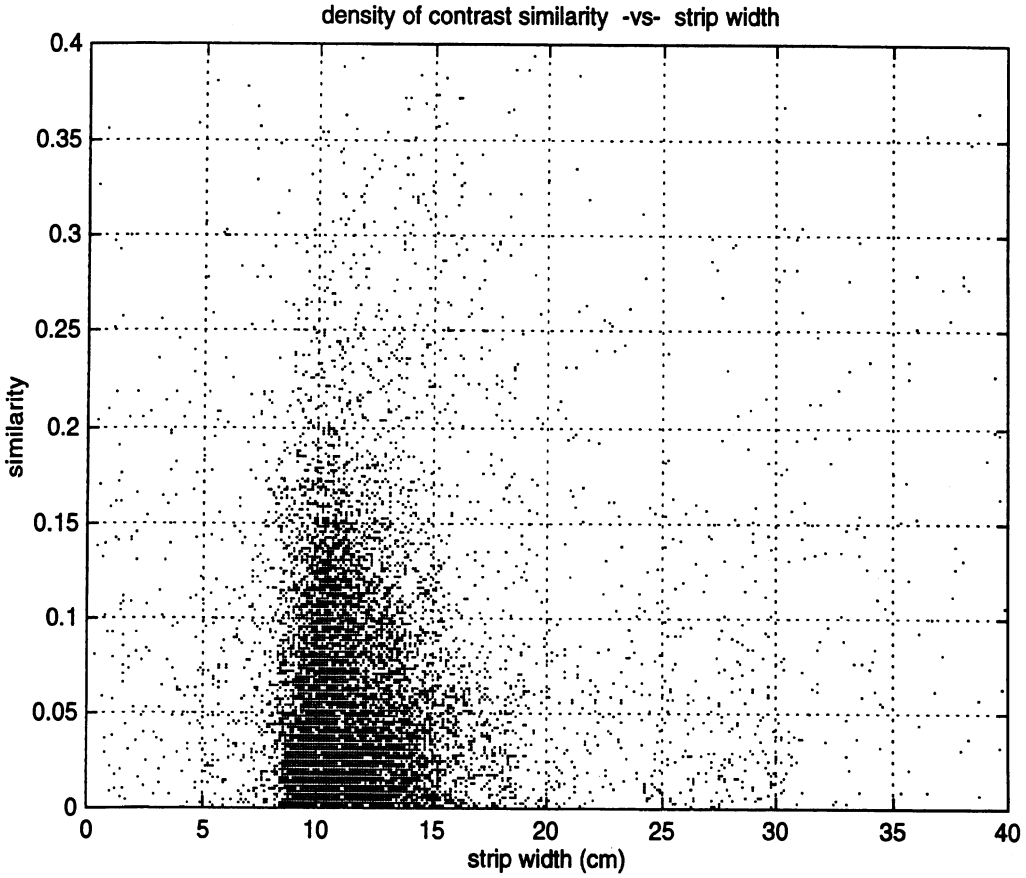

Figure 5.2.1. similarity -vs- stripe width in cm

It can be seen from figure 5.2.1 that two independent thresholds will define a rectangular region (i.e. { $0 \leq \alpha \leq 0.25$, $8 \leq W \leq 15$ }) but the highest density is defined by a triangular or trapezoidal region for this data set. Thus, a threshold that is a linear or higher order combination of width and similarity may provide fewer false alarms than a rectangular decision space.

6.0 CONCLUSIONS

This study produced an experimental test platform and a data set that characterizes some of the aspects of painted-stripe markings as they exist on Michigan highways. The sensing system consisted of active illumination sources, CCD sensors, and analog video recording tape. Data processing was performed on digitized images to locate lane markers and to quantify various properties, such as stripe width and contrast. Example results have been presented for the analysis of over 12000 image fields, and relationships to existing road follower algorithms were highlighted.

The imaging system and data collection platform described in Sections 2-3 allowed for efficient data collection and processing as bounded by the scope of this project. The data processing techniques will correctly extract the edges of the white stripe (or recognize that there is no white stripe) in approximately 95 percent of images. The calibration between image intensity and surface reflectance was only partially completed due to difficulties in correcting for effects involving the illumination geometry, and would need additional work and possible modification of the illumination scheme. Processing time may become an issue with more advanced algorithms, or as the need arises for real-time processing on-board the test vehicle. Attempts to correctly locate the edges of the 5 percent of stripes that the current image segmentation algorithm fails on will require adding complexity to the image processing.

The results presented in Section 5 are an example of the kind of analyses that can be performed with such a data set. Other algorithm developers will likely generate other decision spaces or threshold techniques, and an evaluation of the failure rate based on a simple distribution of example markers will provide valuable insight into performance and where modifications need to be made. Many other parameters have been generated and can be generated, depending on the needs of the end user.

Obviously, improvements can be made in all regions of the study, most notably in the illumination design and calibration technique. Expansion and improvement of the platforms capabilities will likely yield a more expansive data set. A longitudinal study of markers may indicate a "shuffling" of distributions between stages of weather, wear, and restriping. Thus, the distributions presented above are but a "snapshot," and will vary over time.

7.0 REFERENCES

[1] Behringer, Reinhold. "Road Recognition from Multifocal Vision," in *Proceedings, Intelligent Vehicles `94 Symposium*, pp. 302-307, October 1994.

[2] Dickmanns, Ernst Dieter; and Graefe, Volker. "Applications of Dynamic Monocular Machine Vision." *Machine Vision and Applications*. 1:241-261, 1988.

[3] Dickmanns, Ernst D.; and Mysliwetz, Birger D. "Recursive 3-D Road and Relative Ego-State Recognition." *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 14(2):199-213, 1992.

[4] Efenberger, W.; Ta, Q.-H.; Tsinas, L.; and Graefe, V. "Automatic Recognition of Vehicles Approaching from Behind." In *Proceedings of the Intelligent Vehicles `92 Symposium*, pp. 57-62, June 1992.

[5] Franke, U. "Real Time 3D-Road Modeling for Autonomous Vehicle Guidance." *Seventh Scandanavian Conference on Image Analysis*. August 1991.

[6] Franke, U.; Fritz, H.; and Mehring, S. "Long Distance Driving with the Daimler-Benz Autonomous Vehicle VITA." *PROMETHEUS Workshop*. December 1991.

[7] Grimmer, David; and Lakshmanan, Sridhar. "Finding Straight Edges in Radar Images Using Deformable Templates." In *Proceedings of the 27th Conference on Information Sciences and Systems*. 1993.

[8] Hashimoto, K.; Nakayama, S.; Saito, T.; Oono, N.; Ishida, S.; Unoura, K.; Ishii, J.; and Okada, Y. "An Image-Processing Architecture and a Motion Control Method for an Autonomous Vehicle." In *Proceedings of the Intelligent Vehicles `92 Symposium.*, pp. 213-218, June 1992.

[9] Kenue, Surender K. "LANELOK: Detection of Lane boundaries and Vehicle Tracking Using Image-Processing Techniques -- Part I: Hough-Transform, Region Tracing and Correlation Algorithms." *SPIE Mobile Robots IV*, pp. 221-233, 1989.

[10] Kenue, Surender K. "LANELOK: Detection of Lane boundaries and Vehicle Tracking Using Image-Processing Techniques -- Part II: Template Matching Algorithms." *SPIE Mobile Robots IV*, pp. 234-244, 1989.

[11] Kluge, Karl; and Lakshmanan, Sridhar. "A Deformable Template Approach to Lane Detection," In *Proceedings, Intelligent Vehicles `95 Symposium.*, pp 54-59.

[12] Kluge, Karl. "Extracting Road Curvature and Orientation From Image Edge Points Without Perceptual Grouping Into Features," in *Proceedings of the Intelligent Vehicles `94 Symposium*, pp. 109-114, October 1994.

[13] Kluge, Karl, "YARF: An Open-Ended Framework for Robot Road Following." PhD thesis, Carnegie Mellon University, 1993.

[14] Kluge, Karl; and Thorpe, Charles. "Representation and Recovery of Road Geometry in YARF." In *Proceedings of the Intelligent Vehicles `92 Symposium*, pp. 114-119, June 1992.

[15] Kluge, Karl; and Thorpe, Charles, "Intersection Detection in the YARF Road Following System," *Proceedings, Intelligent Autonomous Systems 3*, pp. 145-154, February 1993.

[16] Kuan, Darwin; Phipps, Gary; and Hsueh, A.-Chuan. "Autonomous Land Vehicle Road Following." *Proceedings First International Conference on Computer Vision*, pp. 557-566, June 1987.

[17] Kushner, Todd R.; and Puri, S. "Progress in Road Intersection Detection for Autonomous Vehicle Navigation." *Mobile Robots II*, pp. 19-24, November 1987.

[18] Liou, Shih-Ping; and Jain, Ramesh. "Road Following Using Vanishing Points." *Computer Vision, Graphics, and Image Processing*. 39:116-130, 1987.

[19] Novak, C.; Shafer, S.; Wilson, R. "Obtaining Accurate Color Images for Machine Vision Research" In *Proceedings, SPIE Conference on Perceiving, Measuring, and Using Color*, February 1990

[20] Ohlander, R.; Price, K.; Reddy, D. Raj. "Picture Segmentation Using a Recursive Region Splitting Method." *Computer Graphics and Image Processing* 8(3):313-333. Dec. 1978.

[21] Struck, G.; Geisler, J.; Laubenstein, F.; Nagel, H.-H.; and Siegle, G. "Interaction Between Digital Road Map Systems and Trinoccular Autonomous Driving," in *Proceedings, Intelligent Vehicles `93 Symposium*, pp. 461-465, July 1993.

[22] Struck, G.; Geisler, J.; Laubenstein, F.; Nagel, H.-H.; and Siegle, G. "Multi-Camera Vision- Based Autonomous Maneuvering at Road Intersections," in *Proceedings, Intelligent Vehicles `94 Symposium*, pp. 189-194, October 1994.

[23] Suzuki, A.; Yasui, N.; Nakano, N.; and Kaneko, M. "Lane Recognition System for Guiding of Autonomous Vehicle." In *Proceedings of the Intelligent Vehicles '92 Symposium*, pp. 196-201, June 1992.

[24] Thorpe, Charles; Amidi, Omead; Gowdy, Jay; Hebert, Martial; and Pomerleau, Dean, "Integrating Position Measurement and Image Understanding for Autonomous Vehicle Navigation," *Proceedings of the Second International Workshop on High Precision Navigation*, 1991.

[25] Turk, Matthew A.; Morgenthaler, David G.; Gremban, Keith D.; and Marra, Martin. "VITS -- A Vision System for Autonomous Land Vehicle Navigation." *IEEE Transactions on Pattern Analysis and Machine Intelligence*. 10(3), pp. 342-361, May 1988.

8.0 APPENDIX

Statistical Characterization of the Visual Characteristics of

Painted Lane Markings

from *Proceedings, Intelligent Vehicles '95*

Proceedings of the Intelligent Vehicles '95 Symposium, pp. 54-59
September 25-26, 1995, Detroit, USA

Statistical Characterization of the Visual Characteristics of Painted Lane Markings

Karl Kluge
University of Michigan AI Lab
1101 Beal Avenue
Ann Arbor, MI 48109-2110 (USA)
kckluge@eecs.umich.edu

Greg Johnson
Univ. of Mich. Transportation Research Institute
2901 Baxter Road
Ann Arbor, MI 48109-2150 (USA)
dirk@eecs.umich.edu

Abstract

Most vision-based systems for lane detection and tracking use painted lane markings as the visual cues which determine the location of the camera relative to the lane. Almost all of the work that has been done in the area of evaluating the performance of these systems has focused on the accuracy of the recovered lane geometry. Reliability of feature detection as a function of intrinsic marking properties, ambient lighting and weather conditions, and viewing geometry is an equally important aspect of algorithm performance which must be explored if progress is to continue in this area of research. This paper reports a small scale effort to attack one aspect of this problem, the automated characterization of the intrinsic visual properties of white painted lane markings. Images of the right lane marking are taken by a camera mounted in an trailer enclosure towed behind a vehicle, allowing control of the lighting conditions. The intensity histogram of each image is examined to select a threshold which is used to classify each pixel as pavement or stripe. The edges of the white stripe are located using robust estimation and a shared vanishing point constraint. Once the stripe edges are located in an image, stripe properties such as width, brightness, and contrast with the pavement are calculated.

1 Introduction

Vision-based algorithms for lane detection and tracking have been an active area of research over the last decade (see the discussions of related work in [9] and [10]). Very little work, however, has been done in the area of evaluating the performance of such algorithms. The performance evaluation work that has been done has focused on accuracy of recovered geometry rather than reliability of feature extraction. This is a critical gap, as the ability to characterize the reliability of feature extraction in vision-based lane detection and tracking is critical for several reasons:

- to algorithm developers
  - determining the performance limitations of a particular algorithm
  - testing whether proposed "fixes" to existing known problems actually improve algorithm performance
- to people making funding and deployment decisions:
  - determining deployability of an algorithm
  - deciding between competing algorithms
- to people maintaining highways and Automated Highway System (AHS) lanes:
  - determining which sections of lane markings need repainting
  - determining what conditions make use of an AHS lane unsafe

Evaluation of an algorithm's performance as a function of environmental conditions requires the ability to characterize the intrinsic visual properties of the features being tracked by the algorithm. This work focuses on that aspect of performance evaluation. Specifically, we have constructed a hardware platform and algorithms which enable us to measure the intrinsic visual properties of painted white lane markings.

The next section reviews prior work in the area of performance evaluation of vision-based lane detection and tracking. Section 3 describes the computer vision algorithm used to detect the white stripe and presents the results of experiments to estimate the reliability of the segmentation. Section 4 describes the stripe properties we are currently calculating. Section 5 describes our plan to extend this work into a characterization of the performance of some of the feature trackers used in the YARF road following system [10].

2 Related work

Most prior efforts to quantitatively characterize the performance of vision-based lane tracking techniques have focused on parameters related to road geometry. Pomerleau [13][19] manually determined the location of the road center at one meter intervals along a 140 meter stretch of road. He then set up a siphon at the center of the rear of the NAVLAB vehicle so that it would drip water along the vehicle's path. He used this data to compute the mean and variance of the vehicle's deviation from the center of the road, and compared this to the mean and variance of the vehicle's deviation from road center when

driven by a human driver. Behringer, et al, as part of the VaMoRs project [1] compared the lane-tracking system's estimate of road curvature with measured ground truth on a 450 meter test track. Morganthaler, et al, [11] combined road edge points determined by segmenting a color image with range information from a laser range scanner to determine ground truth for the road geometry in a scene, and used this as a basis for comparing the accuracy of the road shape recovered by three algorithms using different terrain models (flat-earth, hill-and-dale, and zero-bank).

The data used to derive estimates of road geometry and vehicle location on the road derive from image segmentation algorithms. The reliability of these algorithms is therefore another key parameter in evaluating system performance. Kluge [10] compared the performance of three local feature detection algorithms on the white and double yellow lines defining the lane edges in a set of 16 images. The algorithms were applied to a fixed set of rows at columns chosen based on curves fit to a small number of hand-selected points along each lane edge.

None of these efforts characterized the road scene data in terms of physical properties of the road surface or ambient weather and lighting conditions. All except the work by Morganthaler, et al, made use of a limited set of ground truth data measured manually. A NIST report on performance evaluation for robotic vehicle technologies [6] suggested using a side-looking camera to determine lane boundary type and location. This information would be used to calculate the variance of the difference between vehicle offset as measured from the lane edge location in the side-looking camera's field of view and the estimate of vehicle offset provided by a lane-tracking system. That variance would then be correlated with scene characteristics such as lane boundary type and ambient lighting conditions (sun location, rain, clouds, shadows). No suggestions were given as to how to characterize the lane boundary or what specific environmental measurements to make.

A number of computer vision systems have been built to characterize pavement cracking [14][7], but we have been unable to locate any similar work in the area of characterizing painted lane markings. The detection algorithms used in existing lane-tracking systems tend to be fairly strongly model-based, using assumptions about expected width, orientation, and contrast to try to reduce or eliminate false positive feature detections. Since the goal of the work reported here is to detect all kinds of stripes -- the good, the bad, and the ugly -- a more data-driven technique was chosen in order to try to capture weaker, less well-defined stripes. The algorithm chosen uses histogram-based thresholding to identify stripe pixels, followed by line-fitting to the edges of the stripe region to locate the stripe edges. The next section presents the segmentation algorithm in detail.

3 Detecting stripe edges in images

The hardware set-up used in this project is shown in Figure 1. A vehicle tows a trailer enclosure which contains lighting, a camera, and a VCR to record the image data. The camera produces grey-level images, using a shutter speed of 1 / 250th of a second to minimize blurring of the images. Processing to extract the stripe location and calculate stripe characteristics is done off-line in the lab.

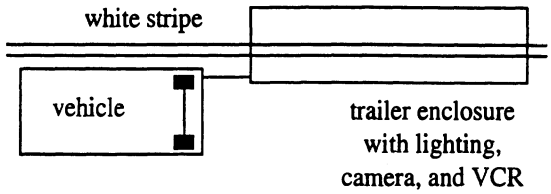

Figure 1: Stripe characterization hardware.

The first stage in the process of characterizing stripe properties is to detect the stripe in the image returned by the camera in the trailer enclosure. The area imaged is fairly small, so only pavement and stripe are imaged (areas corresponding to the enclosure walls can be cropped from the images in software). In addition, the use of the enclosure eliminates illumination variations across the area imaged due to shadows. As a result of the constrained nature of the images a simple segmentation algorithm was chosen which selects an intensity threshold to split the image pixels into pavement (below threshold) and stripe (above threshold) pixels.

The left edge of the stripe is located by finding the left-most stripe pixel in each row and using a robust estimation technique to fit a line to that set of points. The right edge of the stripe is similarly determined by finding the right-most stripe pixel in each row and fitting a line to those points. The stripe edge fits are made subject to the constraint that the left and right stripe edges must converge at a point on the horizon row in the image plane, corresponding to the constraint that the stripe edges are parallel lines on a flat surface.

The thresholding and edge fitting stages of the segmentation algorithm are described in more detail below. After these sections is a description of experimental results used to evaluate the reliability of the algorithm.

3.1. Histogram-based threshold selection

Selection of a threshold based on the characteristics of a histogram of pixel intensities is a classic segmentation technique in computer vision [16]. The type of image dealt with in this work contains a painted white stripe against the fairly uniform background of the dark pavement. As a result, the intensity distribution is close to a bimodal distribution whose peaks correspond to the average pavement and

stripe intensities. The valley between those peaks corresponds to the appropriate threshold to use to classify pixels as pavement or stripe. While the details of the implementation differ, the method used to select the threshold is similar in spirit to the mode method described in [16].

An example is shown in Figure 2. The intensity histogram, shown in the upper right quadrant of the figure, is almost bimodal. There are two valleys in the histogram to the right of the selected threshold, but fairly simple heuristics permit the algorithm to select the correct valley to use. The results of the thresholding are shown in the lower right quadrant of the figure, demonstrating that the threshold selected does a good job separating the image into stripe and pavement pixels.

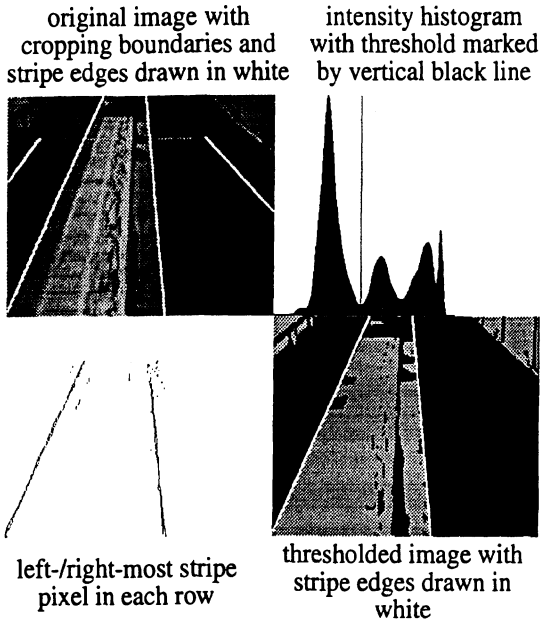

Figure 2: Threshold selection example with bimodal intensity distribution.

In some cases where the stripe is very faint or worn there is no distinct second peak in the intensity histogram. Instead the histogram peak corresponding to the pavement intensities has a long right tail which contains the stripe intensities. In this case the appropriate threshold is selected by finding the intensity value where the pavement peak levels out into the tail of the histogram curve which contains the stripe intensities. An example of this situation is shown in Figure 3.

3.2. Locating the stripe edges

The previous stage of processing, thresholding, classifies each pixel in the image as pavement or stripe. It is still necessary to locate the lines which define the left and right edges of the stripe. This is done in two steps. First, the algorithms locates the left-most and right-most stripe pixel in each row. Second, a line is fit to the set of left-most stripe pixels

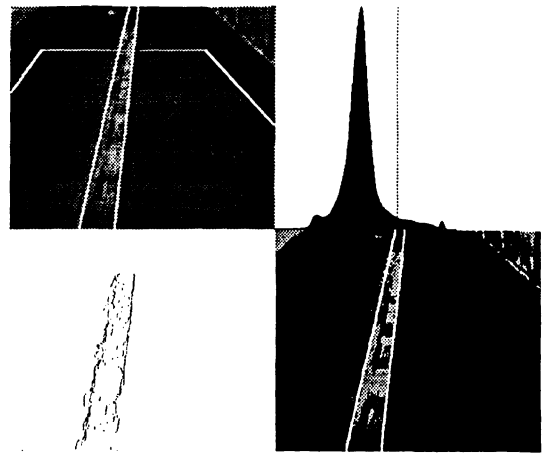

Figure 3: Threshold selection example with long right tail rather than bimodal histogram.

to locate the left stripe edge and a line is fit to the set of right-most stripe pixels to locate the right stripe edge.

This line fitting is complicated by two factors. The first of these factors is that the thresholding process can incorrectly classify bright pavement pixels as stripe rather than pavement. The second of these factors is that wear may result in some pixels inside the stripe edges having an intensity below the threshold level. Two strategies are used to reduce the effects of these outliers on the edge fits.

The first strategy is to reduce the number of degrees of freedom in the edge fits by imposing a constraint on the parameters of the left and right stripe edges. These edges are represented as lines in the image plane of the form $column = m \times row + b$. The area imaged under the enclosure is small, making it reasonable to treat the ground as a flat surface. The short length of stripe visible (3.56 meters) makes it reasonable to approximate the stripe edges as parallel straight lines. Given a pinhole camera model, parallel lines on a flat surface project into the image plane as lines which meet at a point on the horizon (called the vanishing point of the lines). If the camera is not rolled (i.e. the horizon is parallel to the scanlines in the image plane) the image coordinates can be defined so that the horizon row (which may be above the top of the actual camera field of view in the image plane) is row 0. Parallel lines on the ground plane will then have the same column intercept value b.

The second strategy used to increase the robustness of the fit in the presence of outliers in the sets of left- and right-most pixels is to use Least Median of Squares (or LMS) estimation [15] to determine the left and right edge slopes and the shared vanishing point. The number of outlying points in the data set has to equal or exceed 50% before the parameters found by LMS estimation are affected by their presence.

3.3. Segmentation performance evaluation

In order to evaluate the performance of the image segmentation algorithm the program was run on a set of 615 images and the results were hand-classified. Each segmentation was classified as Correct (stripe edges unambiguously correctly located; also, correct identification of images which didn't contain a stripe), Qualified Correct (faint paint adjacent to a located stripe edge makes edge localization arguable), or Failed (incorrect result due to incorrect choice of threshold, false positives, etc.).

Of the 615 images, 499 (81.14%) were rated as Correct. 87 (14.15%) were rated as Qualified Correct. Only 29 (4.72%) were rated as Failures. Typical examples of Qualified Correct results are shown in Figure 4.

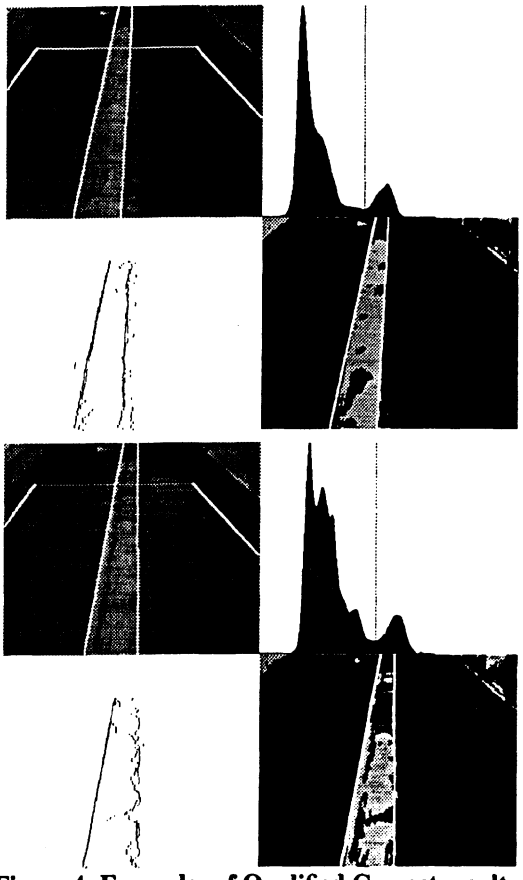

Figure 4: Examples of Qualified Correct results.

The segmentation algorithm takes 1.4 seconds on a SPARC 20 (SPEC floating point rating 78.3) to detect the stripe edges. No effort has been made at this point to optimize the processing to increase the speed of the algorithm. Sampling fewer pixels during histogram construction would be the obvious way to reduce the processing time, but the possible effects of sampling fewer pixels on the reliability of threshold selection would have to be evaluated to insure that performance wasn't degraded.

4 Characterizing visual properties of the white stripe

Once the stripe edges have been detected in an image it is possible to compute measures which characterize the visual properties of the stripe. The current statistics we are computing are:

- stripe width in cm.;
- mean intensity of pixels inside the stripe edges;
- standard deviation of the intensity of pixels inside the stripe edges;
- mean intensity of pixels outside the stripe edge (i.e., pavement pixels);
- standard deviation of the intensity of pavement pixels;
- absolute contrast between the stripe and pavement (defined as the difference in means between the stripe pixel intensities and the pavement pixel intensities); and
- relative contrast between the stripe and pavement (defined as the difference in mean intensities of the stripe and pavement divided by the standard deviation of the stripe intensities).

Currently the intensity-based statistics are uncalibrated. In order to capture the actual intrinsic physical reflectance properties of the stripe and pavement, the raw intensities need to be calibrated to absolute units [5][12]. The enclosure has a calibration card near the top and center of the camera's field of view. We are currently calibrating the camera so that we can use this card to convert raw intensities into absolute albedo values.

Figure 5 shows three plots of stripe width vs. relative contrast for a set of 315 images. These images were taken at approximately 10 second intervals while the vehicle pulling the enclosure trailer drove at normal highway speeds (55 to 65 mph). The top graph shows the values for all images in the data set. This set of images was part of the data set used to evaluate the segmentation algorithm. As a result, it is possible to plot the same data showing only point corresponding to images whose segmentation was classified as Correct (middle graph). The bottom graph shows the data points corresponding to images whose segmentation was classified as Qualified Correct (crosses) or Failure (diamonds) (bottom graph).

As can be seen from a comparison of these plots, Qualified Correct results often occur in images where the stripe is either very worn and degraded (the very narrow stripes) or wide and illdefined due to

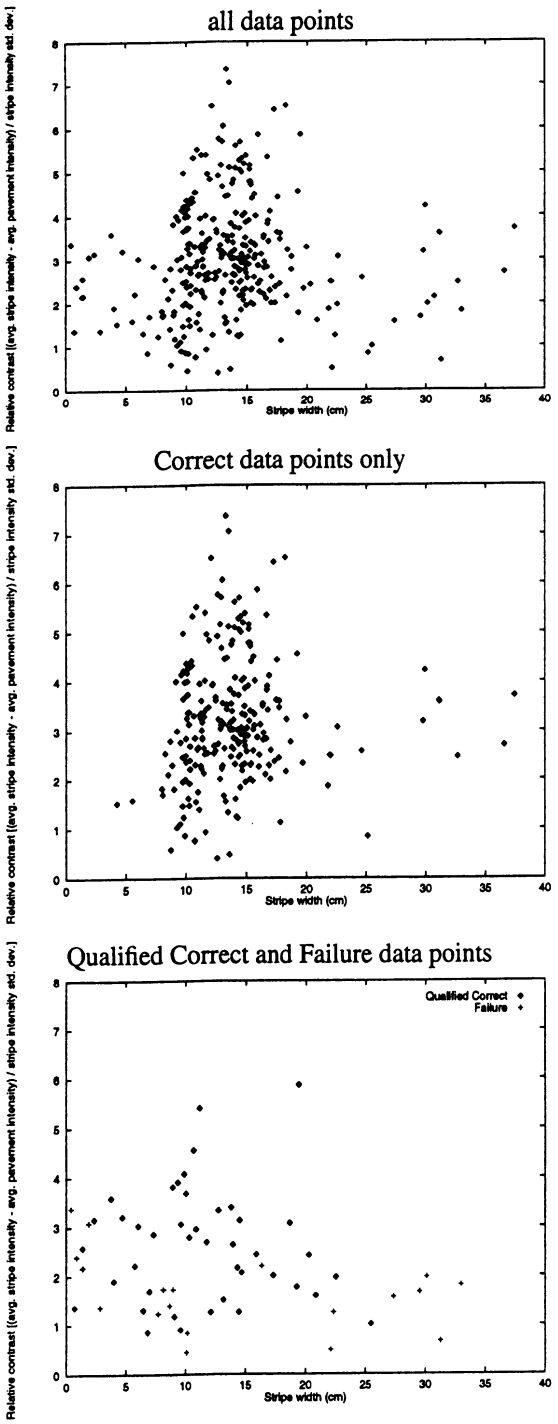

Figure 5: Plots of stripe width vs. relative contrast for a data set of 315 images.

overpainting (the very wide stripes). There doesn't seem to be any clear separation in the stripe parameter space between segmentation failure cases and Qualified Correct cases.

5 Future work

In addition to efforts to improve the reliability and speed of the image segmentation, we plan to extend this effort to characterize the performance of the white stripe tracking algorithms used in the YARF road following system [10]. The ability of an image segmentation algorithm to detect a painted lane marking in the road environment depends on a variety of factors, including

- viewing geometry (distance to the target area of the scene being examined; viewing angle with respect to horizontal to the target area; sun location);
- lighting conditions [in this case, directness of illumination (clear vs. overcast vs. foggy)]; and
- physical properties of the feature in the target area (intrinsic visual properties of the white stripe such as width, contrast, etc. as described above).

The above list is not complete (it does not include amount, rate, and type of precipitation, for instance), but it covers most of the major factors. We hope this research will produce further insight into the issues involved in evaluating the performance of vision-based lane trackers as a function of environmental conditions.

While several other vision-based lane tracking systems ([4] and [3], for example) use local feature detection strategies similar to those used by YARF, other strategies for vision-based lane tracking include neural-net approaches [13], segmentation of the entire image [2] [17], and segmentation everywhere along the expected feature locations [8][18]. Different data analysis strategies will be needed to evaluate such systems, but it should be possible to extend the basic methodology to cover them. The ALVINN neural-net road following system, for example, has two methods for estimating the reliability of the net's output. Those reliability values could be correlated with the environmental measurements.

6 Conclusion

There is a lack of methodologies for evaluating the reliability of visual feature detection as a function of environmental conditions in vision-based lane detection and tracking systems. This is a major barrier to progress in the field (indeed, without such methodologies it is impossible to determine if progress is taking place). We have constructed a system which addresses one aspect of this problem, characterizing the intrinsic visual properties of painted stripes. The image segmentation we are currently using works on 95% of the images examined in a test set, and we intend to work to improve performance further. We also plan to use this system to characterize the performance of some of the feature trackers used in the YARF road follower.

7 Acknowledgments

This work was co-funded by the Great Lakes Center for Truck and Transit Research under contract UM-67 and by the ITS Research Center of Excellence. Both programs provide for sponsorship of transportation research at the University of Michigan.

8 References

[1] Behringer, R.; v. Holt, V.; and Dickmanns, D. "Road and Relative Ego-State Recognition." In *Proceedings of the Intelligent Vehicles '92 Symposium*, June 1992.

[2] Crisman, Jill; and Thorpe, Charles E. "SCARF: A Color Vision System that Tracks Roads and Intersections." *IEEE Transactions on Systems, Man, and Cybernetics* 9(1):49-58, February 1993.

[3] Dickmanns, Ernst D.; and Mysliwetz, Birger D. "Recursive 3-D Road and Relative Ego-State Recognition." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14(2):199-213, 1992.

[4] Franke, U.; Fritz, H.; and Mehring, S. "Long Distance Driving with the Daimler-Benz Autonomous Vehicle VITA." *PROMETHEUS Workshop*, December 1991.

[5] Healey, Glenn E.; and Kondepudy, Raghava. "Radiometric CCD Camera Calibration and Noise Estimation." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 16(3):267-276, March 1994.

[6] Herman, Martin; Szabo, Sandor; Murphy, Karl; Coombs, David; Hong, Tsai-Hong; Scott, Harry; Dagalakis, Nicholas; Goodwin, Kenneth; and Albus, James. *Recommendations for Performance Evaluation of Unmanned Ground Vehicle Technologies*. NIST report NISTIR-5244, August 1993.

[7] Kaseko, Mohamed S.; and Ritchie, Stephen G. "A Neural Network-Based Methodology for Pavement Crack Detection and Classification." *Transportation Research Part C (Emerging Technologies)* 1C(4):275-291, December 1993.

[8] Kenue, Surender K. "LANELOK: Detection of Lane boundaries and Vehicle Tracking Using Image-Processing Techniques -- Parts I and II." In *SPIE Mobile Robots IV*, 1989.

[9] Kluge, Karl. "Extracting Road Curvature and Orientation From Image Edge Points Without Perceptual Grouping Into Features." In *Proceedings of the Intelligent Vehicles '94 Symposium*, October 1994.

[10] Kluge, Karl C. *YARF: An Open-Ended Framework for Robot Road Following*. PhD thesis, Carnegie Mellon University, 1993.

[11] Morgenthaler, David G.; Hennessy, Steven J.; and DeMenthon, Daniel. "Range-Video Fusion and Comparison of Inverse Perspective Algorithms in Static Images." *IEEE Transactions on Systems, Man, and Cybernetics* 20(6):1301-1312, November/ December 1990.

[12] Novak, C.; Shafer, S.; and Wilson, R. "Obtaining Accurate Color Images for Machine Vision Research." In *Proceedings, SPIE Conference on Perceiving, Measuring, and Using Color*, February 1990.

[13] Pomerleau, Dean A. *Neural Network Perception for Mobile Robot Guidance*. PhD thesis, Carnegie Mellon University, 1992.

[14] Ritchie, Stephen G. "Digital Imaging Concepts and Applications in Pavement Management." *Journal of Transportation Engineering* 116(3):287-298, May/ June 1990.

[15] Rousseeuw, Peter J.; and Leroy, Annick M. *Robust Regression and Outlier Detection*. John Wiley & Sons, Inc., 1987.

[16] Sahoo, P. K.; Soltani, S.; and Wong, A. K. C. "A Survey of Thresholding Techniques." *Computer Vision, Graphics, and Image Processing* 41(2):233-260.

[17] Schaaser, L. T.; and Thomas, B. T. "Finding Road Lane Boundaries for Vision Guided Vehicle Navigation." In *Vision-Based Vehicle Guidance*. Ichiro Masaki (editor), Springer-Verlag, 1992.

[18] Schneiderman, H.; and Nashman, M. "Visual Processing for Autonomous Driving." In *Proceedings, IEEE Workshop on Applications of Computer Vision*, 1992.

[19] Thorpe, Charles; Amidi, Omead; Gowdy, Jay; Hebert, Martial; and Pomerleau, Dean. "Integrating Position Measurement and Image Understanding for Autonomous Vehicle Navigation." In *Proceedings of the Second International Workshop on High Precision Navigation*, 1991.