Representation and Recovery of Road Geometry in YARF

Karl Kluge and Charles Thorpe

School of Computer Science, Carnegie Mellon University, Pittsburgh PA 15213

(kck@cs.cmu.edu and cet@cs.cmu.edu)

Abstract

YARF, like several other road following systems, models the road as features which form concentric circular arcs lying in a flat ground plane. This can be thought of as a two-dimensional analog of a generalized cylinder, in which a one-dimensional feature cross section is swept along a spine curve which is a circular arc. Two approximations are made in order to use a linear estimation technique to determine the spine parameters. In the first, the circular spine arc is approximated by a parabola. In the second, data points from different road features are shifted to lie on the spine by translating them parallel to the X axis rather than perpendicular to the (unknown) feature tangents. The first section of the paper demonstrates that errors due to these approximations are reduced substantially by estimating the spine parameters in a data-dependent coordinate system rather than in a vehicle-centered coordinate system.

The second section discusses the problem of estimating the spine parameters in the presence of outlying data. Outliers may be caused by false positive responses from features trackers due to incorrect predictions. They can also arise from unexpected shifts in feature location. The outliers are usually correlated in both these cases. Using standard least squares estimation can result in the vehicle drifting off the road due to errors in the parameter estimates caused by such contaminants in the data. The use of least median squares estimation to overcome these problems is discussed.

1. Introduction

Recovery of road structure from segmentation data poses two issues which must be addressed in designing a road following system: the nature of the representation of road structure and the nature of the process which determines the model parameters given a particular set of segmentation data. Selection of an appropriate road representation and data fitting scheme requires balancing a number of conflicting criteria:

> The accuracy with which the class of models selected can represent the actual structure of the road;
>
> The computational cost of extracting the model parameters from segmentation results; and
>
> The robustness and stability of the fitting process in the presence of noise in the segmentation.

YARF adopts a representation scheme in which the ground is assumed to be locally flat, and the road is modelled as a one dimensional set of features swept perpendicular to a spine curve. The spine curve is approximated locally as a circular arc for computational efficiency. Examination of alternative methods of representation in use suggests that this type of road model is the best currently available for balancing the above criteria.

Analysis of the errors introduced by the approximations made to linearize the circular spine arc model shows that their magnitude depends on the coordinate system in which the arc parameters are estimated. Simulation results are presented to show that the magnitude of the errors introduced by linearizing the circular arc model are small in the range of curvatures of interest when the data is rotated into a "natural" coordinate system before parameter estimation.

YARF incorporates two methods for extracting the spine arc parameters given a set of feature positions on the ground plane: least squares fitting and least median of squares fitting. Least median squares [8] is a robust estimation technique which attempts to eliminate the influence of contaminating data points on the estimate of the model parameters. Such a technique is useful in cases where false positive responses from segmentation algorithms result in outlying data points which would otherwise corrupt the estimate of the model parameters.

2. Discussion of techniques for recovery of road model parameters and methods of road representation

2.1. Methods for recovering model parameters

Three main methods have been used to recover road model parameters given image segmentation data: boundary backprojection, voting in the model parameter space, and statistical fitting techniques.

In boundary backprojection, features detected by the segmentation are backprojected onto the (assumed) ground plane, and consistency constraints are applied to determine which features are part of the road. This is the method used in the VITS [10] , FMC [5] , and U. Bristol [9] systems. Algorithms which recover three dimensional road structure using assumptions of constant road width and zero road bank [2] backproject feature points using assumed image projection geometry. The backprojection process doesn't enforce any higher level constraints on relative feature location, and as a result errors in the image segmentation can produce arbitrary errors in the recovered road shape.

In parameter space voting techniques, detected feature locations vote for all possible roads they are consistent with. This method is used in the SCARF [1] , ALVINN [7] , and U. Michigan [6] algorithms, and in some of the LANELOK [4] algorithms. The main advantage of these techniques is their robustness in the face of large amounts of noise in the segmentation results. The main disadvantage is the difficulty of using voting for models which have more than two or three parameters, resulting in large multidimensional Hough spaces. Also, peak detection in the accumulator space can be difficult.

In statistical fitting procedures, road model parameters are fit using the observed data points and the equations of the road model. Standard techniques such as least squares or robust techniques which are less sensitive to outlying data observations can be used, VaMoRs [3] , YARF, and other of the LANELOK algorithms use this type of technique. Of the available techniques for model parameter recovery, statistical fitting methods have a number of advantages. They are computationally efficient and they have a vast literature of theory, techniques, and tools associated with them.

2.2. Methods for modeling road structure

A variety of schemes have been proposed for representing roads. In order of increasing number of parameters in the model, they are: by steering direction; by linear road segments; by circular arc road segments; by flat road segments with locally parallel edges; and by three dimensional roads constrained to have constant width and no banking.

The simplest road representation is to summarize the segmentation data by a steering direction, independent of the actual road geometry. This is the approach taken in the ALVINN neural net road follower. In principle, ALVINN could learn appropriate steering commands for roads which change slope, bank, etc. In practice, images are backprojected onto a flat ground plane and reprojected from different points of view to expand the range of training images. This may prove to be a limiting factor on hilly roads.

The next simplest road representation is to model the road as linear on a locally flat ground plane (or equivalently, as a triangle in the image plane). The road has three parameters, the road width and two parameters describing the orientation and offset of the vehicle with respect to the centerline of the road. LANELOK and SCARF take this approach. The main limit of this type of scheme is the need to move a sufficiently small distance between road parameter estimates so that the straight path being driven along does not diverge too much from the actual road.

Modeling the road as a cross-section swept along a circular arc explicitly models road curvature but retains the flat earth assumption used in linear models. VaMoRs, YARF, and the U. Bristol system use this approach. The equations describing feature locations can be linearized to allow closed form least squares solutions for the road heading, offset, and curvature, as well as the relative feature offsets.

A more general model of road geometry retains the flat earth assumption, but requires only that road edges be locally parallel, allowing the road to bend arbitrarily. This can be done by projection onto the ground plane (VITS and the FMC system), or in the image plane (work at U. Michigan cited above). The lack of higher order constraint on the road shape can lead to serious errors in the recovered road shape when there are errors in the results of the underlying image segmentation techniques.

Several algorithms have been developed to recover three dimensional variations in road shape under the assumption that the road does not bank [2] . These current algorithms use information from a left and right road edge, which precludes integrating information from multiple road markings. Evaluation of an early zero-bank algorithm by the VITS group as part of the ALV project suggested that such algorithms may be very sensitive to errors in feature location by the segmentation processes. This is due to the assumption of constant road width, which leads to errors in road edge location being interpreted as the result of changes in the terrain shape.

Circular arc models would appear to be the technique of choices in the absence of algorithms for the recovery of three dimensional road structure which are robust in the presence of noise in the segmentation data. They have a small number of parameters, they impose reasonable constraints on the overall road shape, and statistical methods can be used for estimating the shape parameters, with all the statistical theory and tools that use of such methods allows the system to apply to the problem.

2.3. Road model and parameter fitting used in the YARF system

YARF models the road as a one-dimensional feature cross-section swept on a flat ground plane perpendicular to a spine curve (hereafter referred to as a *generalized stripe* model). Such a model lends itself to parameter estimation using statistical fitting techniques and seems to work reasonably well even in the presence of mild variations in ground plane orientation (gentle hills, for instance). Figure 1 shows an image of a two lane divided road. Feature points have been detected along both white lines and the double yellow line in the center. Figure 2 shows the data points and recovered road shape on the ground plane.

YARF assumes that the spine curve can be approximated by a circular arc. In order to have a system which is linear in its parameters a parabolic approximation is made to a circular arc. This parabola represents the binomial series expansion of the circular arc equation. The term representing the displacement of a detected feature point from the spine is also linearized. The final linear model that results from these approximations is $x = curvature \times y^2/2 + heading \times y + spinetrans - offset$, where (x, y) is the position on the ground plane of a detected feature point, *offset* is the offset of the feature from the road spine, *curvature* is the curvature of the spine arc, *heading* is related to the tangent of the spine arc at the x-intercept, and *spinetrans* is the x-intercept of the spine arc.

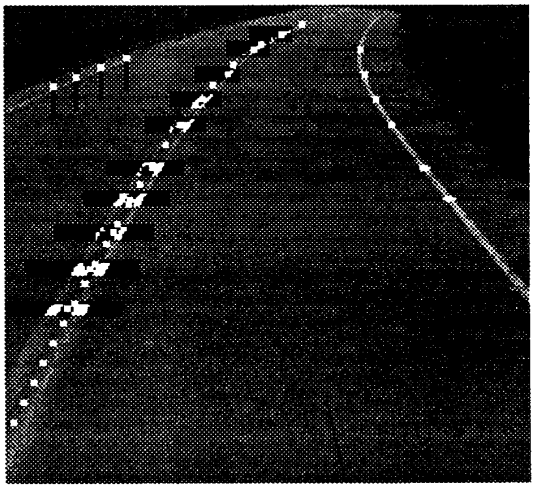

Figure 1: Road image with trackers on lane markings

Standard statistical estimation techniques can be used to recover the spine arc parameters of *curvature*, *heading*, and *spinetrans*. YARF uses either least squares estimation or least median of squares estimation. The next section discuss the errors introduced by the approximations used to derive a linear system, presenting an analysis of the magnitude of the errors introduced. After that, fitting techniques are discussed, with an explanation of least median of squares fitting and why it is preferable in some cases to standard least squares.

3. Errors introduced by linear approximations in YARF

There are two sources of error introduced by the linearizations. The first arises from the approximation of a circular arc by a parabola. The second arises from translating points parallel to the x-axis to move them onto the spine.

3.1. Approximating a circular arc by a parabola

Consider the equation of a half circle centered at the origin, $x = \sqrt{r^2 - y^2}$. This can be expressed as a series, $x = c_0 + c_1 \times y + c_2 \times y^2 + c_3 \times y^3 + \ldots$ Performing the binomial series expansion to solve for the coefficients and ignoring terms beyond y^2 yields the parabola $x = r + y^2/(2r)$. Introducing translation in x simply changes the interpretation of the constant term of the series from $c_0 = r$ to $c_0 = r + x_{center}$. Translation in y makes the coefficient of the y term in the series nonzero by substituting $y' = y - y_{center}$ into the parabola equation above.

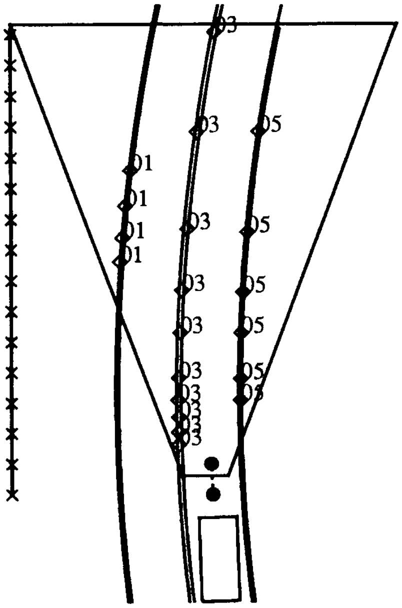

Figure 2: Reconstructed road model

The axis of the parabolic fit is implictly the y value about which the series approximation of the spine arc is being expanded. The further data points lie from that axis, the greater the divergence between the estimated arc parameters and the actual arc parameters. Since the fit constrains the axis of the parabola to be parallel to the x axis, rotating the data so that the x axis passes through the mean y value of the data points reduces the fraction of the arc circumference spanned by the fit, and increases the accuracy of the estimate of the spine arc parameters.

3.2. Translating data points parallel to the X axis rather than perpendicular to the arc

Data points from the features being tracked must be translated to lie on the spine in order to fit the spine parameters. The translation is made parallel to the x axis rather than perpendicular to the (unknown) spine arc in order to keep the problem linear (see Figure 3 below). The magnitude of the error introduced is $error = (x - x_{center}) - offset + \sqrt{((x - x_{center})^2 - offset^2 - (2 \times offset \times radius))}$.

The magnitude of this error is also dependant on the coordinate system chosen for the fit. Again, rotating the data so that the x axis is roughly perpendicular to the predicted road at the mean y value of the data spreads the error more evenly among the points and reduces the size of the error for the points with larger y values.

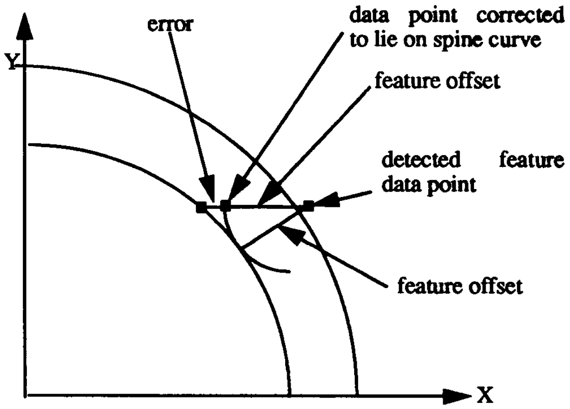

Figure 3: Error introduced by translating points to road spine parallel to the X axis

3.3. Evaluation of error introduced by linear approximations to circular arc model: Simulation results

Simulations were run in order to provide quantitative estimates of the errors introduced by the linearizations described above. These simulations demonstrate the importance of fitting the data in a "natural" coordinate system in which the *x* axis is perpendicular to the road at the mean y value of the data points. We call such rotation of the data before parameter estimation *virtual panning*. The camera and road geometry models from an actual Navlab run were used to generate synthetic road images of specified curvature, and YARF was run on the synthetic images to gather data on the difference between the estimated road shape and the actual road shape.

The simulated vehicle drove 2 meters between images, keeping centered in the lane with the rear axle perpendicular to the spine curve. The simulator was set up to use the same road and camera models to generate the images and to backproject and fit the data, and the image data is idealized. This eliminates sources of error other than the approximations described above. After allowing the simulation to run for 10 frames to allow the system to settle into a steady state, the fits from the eleventh to twentieth frames were averaged and compared to the known model.

The error measure chosen was distance from the true lane center to the estimated lane center at a given distance along the estimated lane center arc. The error was plotted for distances along the estimated lane center starting at the rear axle of the vehicle and extending out to 40 meters. The front end of the vehicle is about 3.5 meters in front of the rear axle. The error measure is illustrated in Figure 4 below.

Figure 5 compares the error in the estimated lane center position with and without virtual panning of the data. The top graph shows the error for positive radii of curvature with the estimation done in a coordinate system fixed with respect to the vehicle. Note that in all cases the error in the range of 5 to 15 meters from the rear axle is very small, staying within a foot of the actual lane center. As the curves become tighter the errors increase dramatically for distances greater than 20 meters along the estimated lane center.

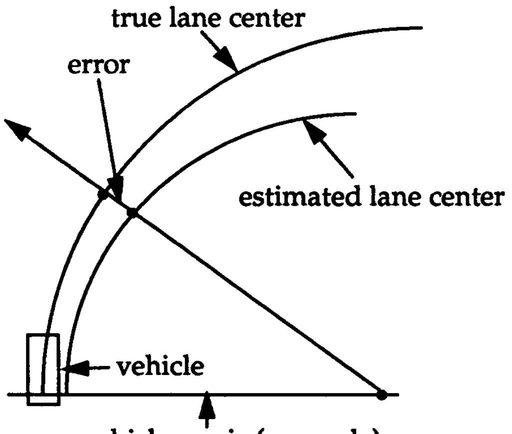

Figure 4: Illustration of estimation error measure for simulations

The bottom graph in Figure 5 shows the corresponding errors when the data is virtually panned prior to estimation of the spine parameters. The vertical scale is not the same in this graph in order to improve readability. Note that the magnitude of the error stays under 80 cm. at all distances out to 40 meters along the estimated lane center, and for all radii of curvature down to +/- 30 meters. This shows the improvement in fit accuracy achieved by rotating the data into a "natural" coordinate system. Implementation of virtual panning has permitted the YARF system to successfully navigate curves which it could not track previously due to errors in feature prediction.

4. Parameter estimation by Least Median of Squares fitting

Data can be contaminated by observations which do not come from the process whose parameters are being estimated. Such observations are called *outliers*. Their presence in a data set can result in parameter fits that are grossly incorrect when standard least squares techniques are used as estimators. Outliers pose a particular problem for the YARF system. They will arise when there is a false positive response from a tracker. Because the tracker windows are placed at the predicted road position, they will not be random and may pull the fit incorrectly towards the prediction and away from the actual road. In addition, unexpected shifts in feature location will produce data points which are correct responses from the feature trackers, but whose positions are not consistent with the model of the road geometry. An example of this can be seen as an exit ramp is approached. The right edge of the lane veers off, while the left edge continues.

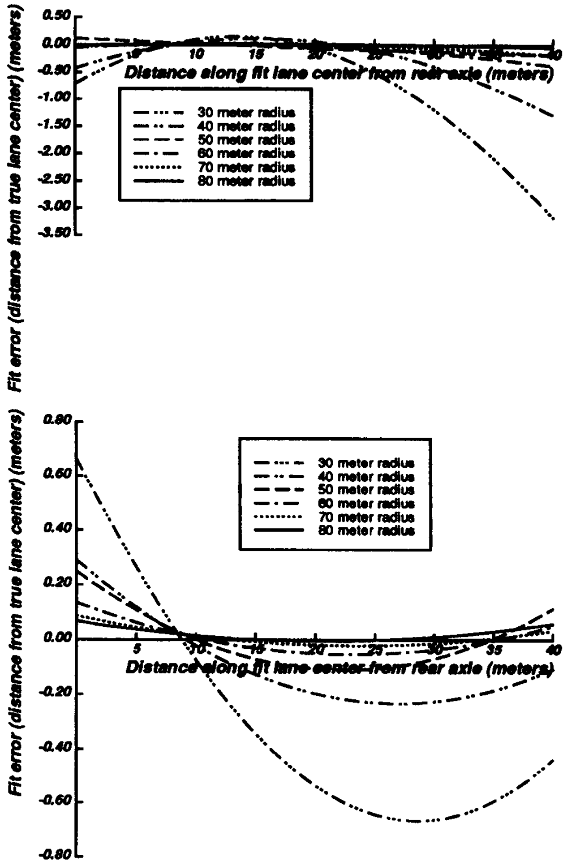

Figure 5: Comparison of accuracy of lane center estimate, fit done in vehicle frame (top) and fit done with virtual panning (bottom)

A number of techniques have been developed to estimate parameters reliably in the presence of outliers. An increasingly popular robust estimation techniques is called *Least Median of Squares* (or LMS) estimation [8]. Consider the linear system $y_i = \beta x_i + \varepsilon$, where β is the vector of parameters to be estimated, and ε is a noise term. The least median squares estimate of β is the $\hat{\beta}$ which minimizes $median(r_i^2)$, where r_i is the residual of the i^{th} data point, $\hat{\beta}x_i - y_i$. To give a simple geometric intuition for what the LMS estimate is, picture the two dimensional linear case. The LMS estimate is the line such that a band centered on the line which contains half the data points has the minimum height in y (the dependent variable) (see Figure 6).

The computation of the LMS estimate is straightforward. Random subsets of the data are chosen. The least squares estimate of the parameters is made for each subset, and the median squared residual for that estimate is computed. The estimate which produced the lowest median squared residual is selected as the final estimate. The LMS estimator can accommodate up to 50% of the data consisting of contaminants, and is less sensitive to correlated outliers than other techniques.

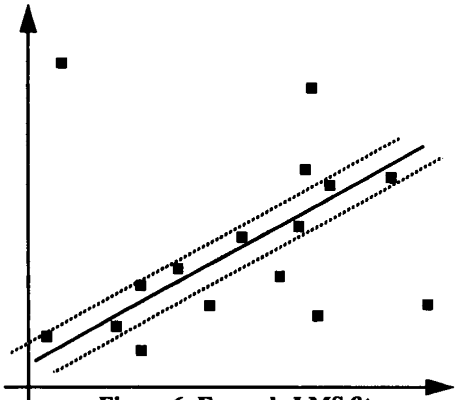

Figure 6: Example LMS fit

Figure 7 shows a case in which the outliers are the result of errors in the segmentation. The lane being followed has a double yellow line on the left side and a single solid white line on the right side. Due to error in the predicted road location some of the trackers for the double yellow line are off the road on grass and return false feature locations. The road actually curves off to the right, explaining the tendency of the points from the left lane edge to fall to the right of the fit in the middle of the diagram, and the failure of the white stripe tracker to locate the right lane edge in that same area (the points marked with asterisks). The LMS fit on the right shows the correct road fit, with the erroneous feature points at top far off to the left of the lane. LMS fitting has been implemented in the YARF system, and has been used successfully to estimate the road parameters during runs on our test site.

5. Conclusion

In this paper we have explained the motivation behind YARF's selection of road representation and model fitting algorithms. We have shown how an analysis of the errors introduced by linearizing the circular arc road model leads to the idea of performing virtual panning on the data to reduce the errors in the model parameter estimates, and presented quantitative results from simulation runs to show the improvement from virtual panning. Also, we have explained the motivation for using least median squares estimation to avoid errors caused by outlying data points

Future work will involve an attempt to characterize the errors induced in the parameter estimates in cases where the flat ground plane assumption doesn't hold. Also there is a need for the development of algorithms which can recover three dimensional road structure with less sensitivity to noise than current algorithms, and which can incorporate information from road features other than the edges of the lane being followed.

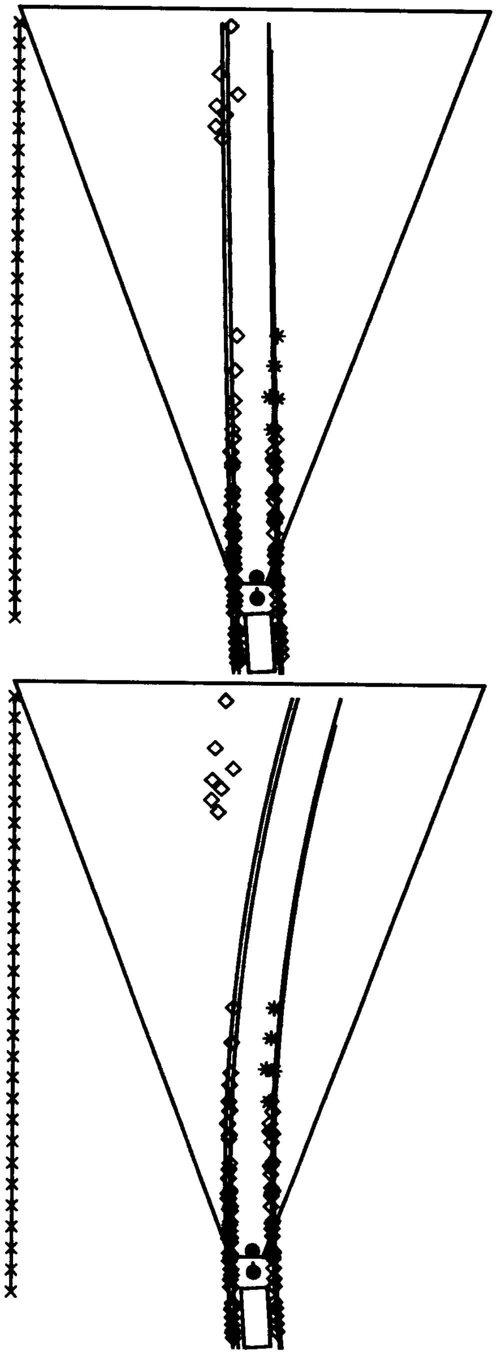

Figure 7: Comparison of least squares fit (left) and least median squares fit (right) of data with outliers.

6. Acknowledgements

This research is partly sponsored by contracts from DARPA, titled "Robot System Development and Testing" (monitored by TACOM) and "Perception for Outdoor Navigation" (monitored by TEC). Our thanks also to the other members of the CMU Unmanned Ground Vehicles group.

7. References

[1] Crisman, Jill. *Color Vision for the Detection of Unstructured Roads and Intersections.* PhD thesis, Carnegie-Mellon University, 1990.

[2] DeMenthon, Daniel, and Davis, Larry. Reconstruction of a Road by Local Image Matches and Global 3D Optimization. *Proceedings 1990 IEEE International Conference on Robotics and Automation.* May, 1990.

[3] Dickmanns, Ernst Dieter and Graefe, Volker. Applications of Dynamic Monocular Machine Vision. *Machine Vision and Applications.* 1241-261, 1988.

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

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

[6] Polk, Amy, and Jain, Ramesh. A Parallel Architecture for Curvature-Based Road Scene Classification. *Roundtable Discussion on Vision-Based Vehicle Guidance '90 (in conjunction with IROS).* July, 1990.

[7] Pomerleau, Dean A. Neural Network Based Autonomous Navigation. *Vision and Navigation: The Carnegie Mellon Navlab.* In Charles E. Thorpe, Kluwer Academic Publishers, 1990, Chapter 5.

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

[9] Schaaser, L. T., and Thomas, B. T. Finding Road Lane Boundaries for Vision Guided Vehicle Navigation. *Roundtable Discussion on Vision-Based Vehicle Guidance '90 (in conjunction with IROS).* July, 1990.

[10] 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), May, 1988.