Title

Development of Binocular Stereopsis for Vehicle Lateral Control, Longitudinal Control and Obstacle Detection

Authors

Malik, Jitendra
Kosecka, Jana
Taylor, Camillo J.
et al.

Publication Date

1999-11-01

CALIFORNIA PATH PROGRAM
INSTITUTE OF TRANSPORTATION STUDIES
UNIVERSITY OF CALIFORNIA, BERKELEY

Development of Binocular Stereopsis for Vehicle Lateral Control, Longitudinal Control and Obstacle Detection

**Jitendra Malik, Jana Kosecka,
Camillo J. Taylor, Philip McLauchlan**

California PATH Research Report
UCB-ITS-PRR-99-37

This work was performed as part of the California PATH Program of the University of California, in cooperation with the State of California Business, Transportation, and Housing Agency, Department of Transportation; and the United States Department of Transportation, Federal Highway Administration.

Report for MOU 306

November 1999

ISSN 1055-1425

This paper uses Postscript Type 3 fonts. Although reading in on the screen is difficult it will print out just fine.

Development of Binocular Stereopsis for Vehicle Lateral Control, Longitudinal Control and Obstacle Detection

PATH MOU-306 Final Report

Department of Electrical Engineering and Computer Sciences
University of California at Berkeley
Berkeley, CA 94720

Principal Investigator: Prof. Jitendra Malik
Postdoctoral Researches:
Jana Košecká
Camillo J. Taylor
Philip McLauchlan

Contents

Abstract

This final report describes the application of computer vision techniques to the lateral and longitudinal control of an autonomous highway vehicle.

In the part of the project we focused on an analysis of the vehicle's lateral dynamics and the design of an appropriate controller for lateral control and investigated various static feedback strategies where the measurements obtained from vision, namely offset from the centerline and angle between the road tangent and the orientation of the vehicle at some look-ahead distance, are directly used for control. The role of the look-ahead, its relation to the vision processing delay, longitudinal velocity and road geometry was crucial on the design of the control and their experimental evaluation.

We carried out a thorough analysis of the effects of changing various important system parameters like the vehicle velocity, the lookahead range of the vision sensor and the processing delay associated with the perception and control systems. We also present the results of a series of experiments that were designed to provide a systematic comparison of a number of control strategies. The control strategies that were explored include a lead-lag control law, a full-state linear controller and input-output linearizing control law. Each of these control strategies was implemented and tested at highway speeds on our experimental vehicle platform, a Honda Accord LX sedan.

For the longitudinal control problem, we investigated the possibility of using stereo vision to provide the range information, in conjunction with a scanning laser radar sensor. The vision based tracking system utilizes a layered architecture wherein the bottom layer computes motion in both images using a simple correlation algorithm, and the upper level performs stereo fixation and reconstruction using an algorithm designed for active vision systems. We present some initial results comparing the quality of range measurements provided by a vision system with the laser radar system.

We report the results from the experimental demonstration of the system as part of the National Automated Highway Systems Consortium (NAHSC) Demonstration which took place in August 1997 in San Diego. The overall system was demonstrated as a part of the main highway scenario as well as part of a small public demonstration of the vision based lateral control on a highly curved test track.

1 Introduction

This report describes the work and research results carried out under MOU-306. Within this proposal we explored the feasibility of the use of visual sensing as a part of the Advanced Vehicle Control System (AVCS). The basic theoretical foundation on which this work was based had been developed under a previous program [39]. In this phase of the research program we demonstrated improved versions of the proposed algorithms, real-time implementations and a novel stereo tracking algorithm for longitudinal control. The integration of the vision subsystem with the vehicle control subsystem enabled us to perform real-world experiments with vision as an integral component of the vehicles control system.

The experimental part of the work has been carried out in collaboration with Honda R&D North America Inc. and Honda R&D Company Limited who provided us with three Honda Accords and a team of Honda engineers responsible for the low-level control of steering and throttle actuators as well as maintenance and overall organization leading towards the final demonstration which was part of NAHSC DEMO 1997 in San Diego.

Following the NAHSC Demo we carried out thourough analysis of the data aquired during numerous experiments and compared the performance of the various control strategies. The control strategies that were examined include a lead lag control law, a full-state linear controller and input-output linearizing control law. Each of these control strategies was implemented and tested at highway speeds on our experimental vehicle platform, a Honda Accord LX sedan. These experiments allowed us to verify the accuracy and efficacy of our modeling and control techniques.

The problem of steering a car along a curved road can be divided into two parts: sensing and control. The *sensing* part involves the extraction of relevant features in the time-varying images and the *control* part deals with the design of the steering control law. Different aspects of steering problem have been examined in the past, both in the engineering and the psychophysics literature.

Several sensing technologies have been proposed for use in an Advanced Vehicle Control System, including vision, magnetic sensors and active range sensors. Some of the most influential work on visually guided control of autonomous vehicles has been done by E.D.Dickmanns and his colleagues [26]. In their system vision was used to provide input for both lateral and longitudinal control of the vehicle on free roads as well as in the presence of other vehicles. The basic approach was to recursively estimate a set of road and vehicle state parameters which included the horizontal and vertical curvature of the road, the vehicle heading angle, the slip angle and the lateral offset of the vehicle with respect to the road. A dynamical model of the vehicle captured their knowledge

about the motion of the vehicle and served as a tool for both fusion of sensor data, design of control strategies and prediction of the effects of control input on the evolution of the measured data.

In the United States one of the leading efforts has been the CMU NavLab Project [36] which has employed a variety of different sensing and control strategies including a neural network based lane following algorithm called ALVINN. A number of other groups throughout the country explored the possibility of using visual sensing for vehicle guidance both in outdoor and indoor environments [37, 33] concentrating primarily on the lane keeping or path following modality on the free roads.

For vision-based lateral control we undertook to explore a strategy which uses direcly the information from the vision at some look-ahead distance. Taking into account the vehicle dynamic model, we formulated the vision dynamics as well as the control objective at the look-ahead distance. This resulted in the development of simpler control strategies Within this setting we evaluated a variety of feedback strategies including lead-lag control, state feedback via pole-placement and I/O linearization. We also implemented observer schemes which provided estimates for the curvature of the roadway which allowed us to incorporate a feedforward component into our control laws. This feedforward component was particularly crucial on highly curved roads.

In case of vision-based algorithms for longitudinal control, in spite of the fact that the context of highways and vehicles is clearly very structured, we avoided using direct scene models in the low-level tracking algorithms, and this distinguishes our work from that of Dickmanns's group [5], for instance. We draw on the large amount of work on scene reconstruction from multiple images in unstructured scenes, in particular the work on robust motion segmentation [22], and affine reconstruction [20]. These approaches are able to take advantage of the redundant information in images, because they latch onto whatever features are available, whereas model-based methods are restricted to the features associated with the chosen model. Redundancy is a vital issue here, because vision is a massively redundant sensor, and approaches which negate this aspect are likely to be discarded in the long term. Because we reconstruct the geometry of the lead vehicle, the algorithms generalize naturally to other vehicle types.

Section 1 presents the basic equations that we have used to model the dynamics of our vehicle and our sensing system and discusses some of the consequences of this model. Section 2 describes the strategy used to extract lane markings from the video imagery and section 3 describes the design of an observer that we use to estimate the states of our system and the curvature of the roadway. Section 4 describes the various control strategies that we implemented on our experimental platform

and section 5 presents the results of the experiments that we carried out with these controllers. Section 6 contains description of a vision subsystem for longitudal control and obstacle detection. Section 7 contains the conclusions that we have drawn from these experiments. A brief description of our experimental vehicle is provided in section 8.

1.1 Vehicle Dynamics

The dynamics of a passenger vehicle can be described by a detailed 6-DOF nonlinear model [34]. Since it is possible to decouple the longitudinal and lateral dynamics, a linearized model of the lateral vehicle dynamics is used for controller design. The linearized model of the vehicle retains only lateral and yaw dynamics, assumes small steering angles and a linear tire model, and is parameterized by the current longitudinal velocity. Coupling the two front wheels and two rear wheels together, the resulting bicycle model (Figure 1.1) is described by the following variables and parameters:

$\boldsymbol{v}$ linear velocity vector (v_x, v_y), v_x denotes speed

α_f, α_r side slip angles of the front and rear tires

$\dot{\psi}$ yaw rate

δ_f front wheel steering angle

δ commanded steering angle

m total mass of the vehicle

I_ψ total inertia vehicle around center of gravity (CG)

l_f, l_r distance of the front and rear axles from the CG

l distance between the front and the rear axle $l_f + l_r$

c_f, c_r cornering stiffness of the front and rear tires.

A simple linear model captures the interaction of the tires with the road surface as follows:

$$\begin{aligned} F_f &= c_f \alpha_f \\ F_r &= c_r \alpha_r \end{aligned} \tag{1}$$

where the side slip angles α_f and α_r between the steering angle and the tire velocity can be expressed as functions of the vehicles kinematic parameters:

$$\begin{aligned} \alpha_f &= \delta_f - \arctan\left(\frac{v_y + l_f\dot{\psi}}{v_x}\right) \approx \delta_f - \frac{v_y + l_f\dot{\psi}}{v_x} \\ \alpha_r &= -\arctan\left(\frac{v_y - l_r\dot{\psi}}{v_x}\right) \approx \frac{-v_y + l_r\dot{\psi}}{v_x} \end{aligned} \tag{2}$$

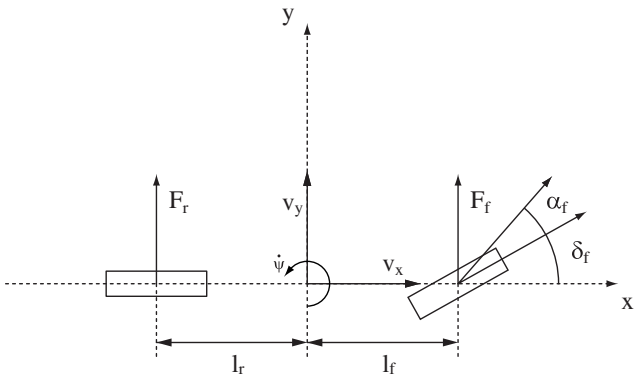

Figure 1: The motion of the vehicle is characterized by its velocity $\boldsymbol{v} = (v_x, v_y)$ expressed in the vehicle's inertial frame of reference and its yaw rate $\dot{\psi}$. The forces acting on the front and rear wheels are F_f and F_r, respectively.

Following Newton law's the net lateral force F and the net torque τ at the center of the gravity are:

$$\begin{aligned} F = F_f + F_r &= ma = m(\dot{v_y} + v_x\dot{\psi}) \\ \tau = F_f l_f - F_r l_r &= I_\psi \ddot{\psi} \end{aligned} \tag{3}$$

Choosing $\dot{\psi}$ and v_y as state variables the lateral dynamics of the vehicle have the following form:

$$\begin{bmatrix} \dot{v}_y \\ \\ \ddot{\psi} \end{bmatrix} = \begin{bmatrix} -\frac{c_f+c_r}{m v_x} & \frac{c_r l_r - c_f l_f}{m v_x} - v_x \\ \\ \frac{-l_f c_f + l_r c_r}{I_\psi v_x} & -\frac{{l_f}^2 c_f + {l_r}^2 c_r}{I_\psi v_x} \end{bmatrix} \begin{bmatrix} v_y \\ \\ \dot{\psi} \end{bmatrix} + \begin{bmatrix} \frac{c_f}{m} \\ \\ \frac{l_f c_f}{I_\psi} \end{bmatrix} \delta_f \tag{4}$$

This linear model for the lateral dynamics and yaw rate is usually referred to as the bicycle model.

1.2 Vision Dynamics

The additional measurements provided by the vision system (see Figure 2) are:

y_L the offset from the centerline at the lookahead,

ε_L the angle between the tangent to the road and the vehicle orientation

Where L denotes the lookahead distance of the vision system as shown in Figure 2. The equations capturing the evolution of these measurements due to the motion of the car and changes in the road geometry are:

$$\dot{y}_L = v_x\, \varepsilon_L - v_y - \dot{\psi} L \tag{5}$$

$$\dot{\varepsilon}_L = v_x\, K_L - \dot{\psi} \tag{6}$$

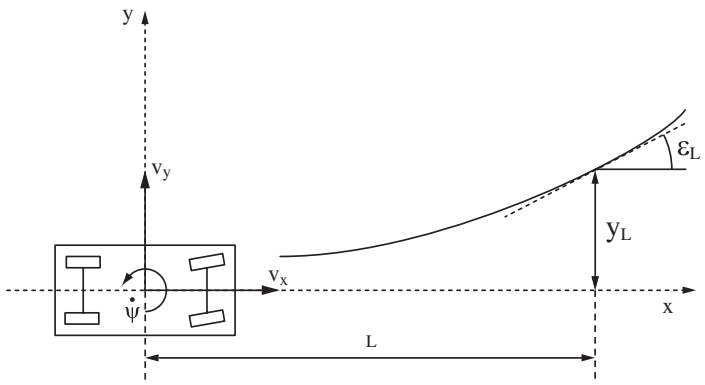

Figure 2: The vision system estimates the offset from the centerline y_L and the angle between the road tangent and heading of the vehicle ε_L at some lookahead distance L.

Where K_L represents the curvature of the road.

1.3 Combined Model

We can combine the vehicle lateral dynamics and the vision dynamics into a single dynamical system of the form:

$$\begin{aligned} \dot{\boldsymbol{x}} &= A\,\boldsymbol{x} + B\,\boldsymbol{u} + E\,\boldsymbol{w} \\ \boldsymbol{y} &= C\,\boldsymbol{x} \end{aligned}$$

with state vector $\boldsymbol{x} = [v_y, \dot{\psi}, y_L, \varepsilon_L]^T$, output $y = [\dot{\psi}, y_L, \varepsilon_L]^T$ and control input $\boldsymbol{u} = \delta_f$. The road curvature K_L enters the model as an exogenous disturbance signal $\boldsymbol{w} = K_L$.

The resulting combined model is captured in Equations (7) and (8).

$$\begin{bmatrix} \dot{v}_y \\ \ddot{\psi} \\ \dot{y}_L \\ \dot{\varepsilon}_L \end{bmatrix} = \begin{bmatrix} -\frac{c_f+c_r}{m v_x} & -v_x + \frac{c_r l_r - c_f l_f}{m v_x} & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \frac{-l_f c_f + l_r c_r}{I_\psi v_x} & -\frac{{l_f}^2 c_f + {l_r}^2 c_r}{I_\psi v_x} & 0 & 0 \\ -1 & -L & 0 & v_x \\ 0 & -1 & 0 & 0 \end{bmatrix} \begin{bmatrix} v_y \\ \dot{\psi} \\ y_L \\ \varepsilon_L \end{bmatrix} + \begin{bmatrix} \frac{c_f}{m} \\ \frac{l_f c_f}{I_\psi} \\ 0 \\ 0 \end{bmatrix} \delta_f + \begin{bmatrix} 0 \\ 0 \\ 0 \\ v_x \end{bmatrix} K_L \qquad (7)$$

$$y = \begin{bmatrix} & & & \\ 0 & 1 & 0 & 0 \\ & & & \\ 0 & 0 & 1 & 0 \\ & & & \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} v_y \\ \\ \dot{\psi} \\ \\ y_L \\ \\ \varepsilon_L \end{bmatrix} + \begin{bmatrix} 0 \\ \\ 0 \\ \\ 0 \end{bmatrix} \delta_f \tag{8}$$

1.4 Analysis

The goal of our analysis is to understand how the behavior of the vehicle varies as a function of important system parameters. In order to do this, we will consider the transfer function $V(s)$ between the steering angle, δ_f, and the offset at the lookahead, y_L. This transfer function can be obtained from the state equations in the usual manner and has the following form:

$$V(\boldsymbol{s}) = \frac{y_L(\boldsymbol{s})}{\delta_f(\boldsymbol{s})} = \frac{1}{\boldsymbol{s}^2} \frac{n_1\boldsymbol{s}^2 + n_2\boldsymbol{s} + n_3}{d_1\boldsymbol{s}^2 + d_2\boldsymbol{s} + d_3} \tag{9}$$

Notice that the transfer function has a pair of poles fixed at the origin along with two poles and two zeros which characterize the dynamic behavior of the vehicle. The coefficients of the denominator of this expression, and hence, the poles of the system depend upon the vehicle velocity v_x. While the numerator terms depend on both the vehicle velocity and the lookahead distance L.

Velocity Figure 3a shows the root locus of the transfer function $V(s)$ for various values of the vehicle velocity v_x assuming a fixed lookahead distance, L, of 10 meters. As the velocity is increased from 10 m/s to 30m/s the poles and the zeros of the transfer function move towards the right half plane and the system becomes less stable.

Lookahead Figure 3b shows how the zeros of the transfer function $V(s)$ are affected by changes in the lookahead distance L. As the lookahead distance is increased the zeros of the transfer function move closer to the real axis which improves their damping ratios. The poles of the transfer function are unaffected by L since this parameter only appears in the numerator of $V(s)$.

Delay Another parameter which affects the behavior of the closed loop lateral control system is the latency associated with the vision system. This can be modeled as a pure delay element with transfer function $D(\boldsymbol{s}) = e^{-T_d\boldsymbol{s}}$ which is placed in series with the vehicles transfer function $V(s)$. The processing delay T_d in our implementation was 57 milliseconds.

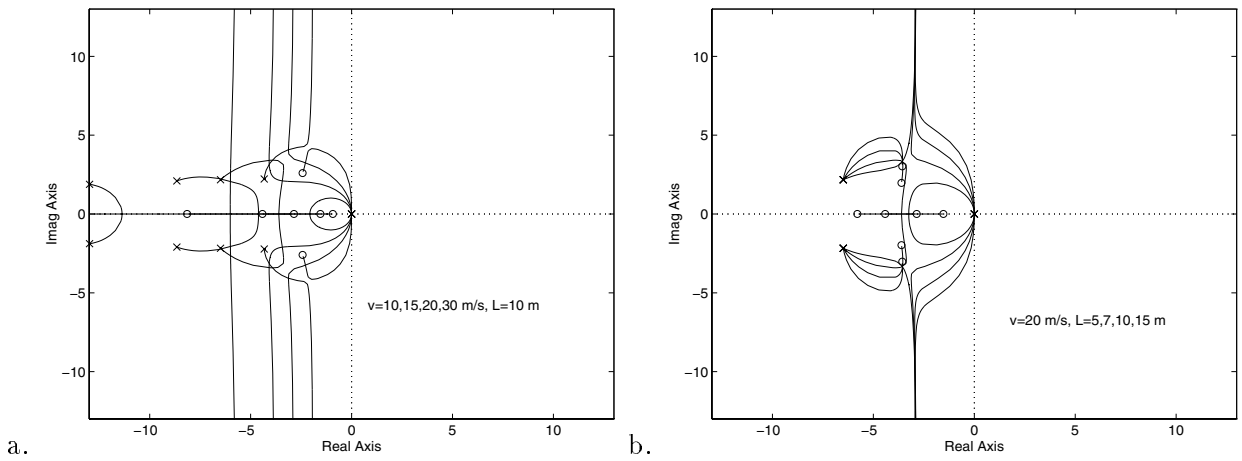

Figure 3: (a) Root locus of $V(\boldsymbol{s})$ for velocities v_x = 10,15,20,30m/s and fixed look-ahead distance L = 10m. Increasing the velocity v_x moves both the poles and zeros towards the imaginary axis. (b) Increasing the look-ahead distance L moves the zeros of the transfer function closer to the real axis, which improves their damping. Once they reach the real axis, further increasing of look-ahead doesn't have any effect on damping. The poles of the transfer function are not affected by changes in L since the parameter appears only in the numerator of $V(\boldsymbol{s})$.

The interplay between the lookahead distance L and the processing delay T_d can be demonstrated quite effectively in the frequency domain. Ideally the overall system should have infinite gain margin and about 40-60° phase margin at the crossover frequency. Bode diagrams of $V(\boldsymbol{s})$ and $V(\boldsymbol{s})\ D(\boldsymbol{s})$ in Figure 4 demonstrate the effect of look-ahead both in the absence (a) and presence (b) of the delay.

Increasing the look-ahead distance adds substantial phase lead at the cross-over frequency. In the presence of processing delay the look-ahead is still able to provide non-zero phase margin for the combined system. For a particular setting of v = 20m/s and L = 20m, the maximum processing delay one can afford to tolerate before bringing the phase margin to zero is $T_{dmax} = 0.39$s, at slower velocities the maximum allowable delay becomes larger. Since the delay adds an additional phase lag over the whole range of frequencies the system bandwidth is clearly limited. From this analysis we can conclude that the delay in our system can be compensated by the additional phase lead provided by increasing the look-ahead distance.

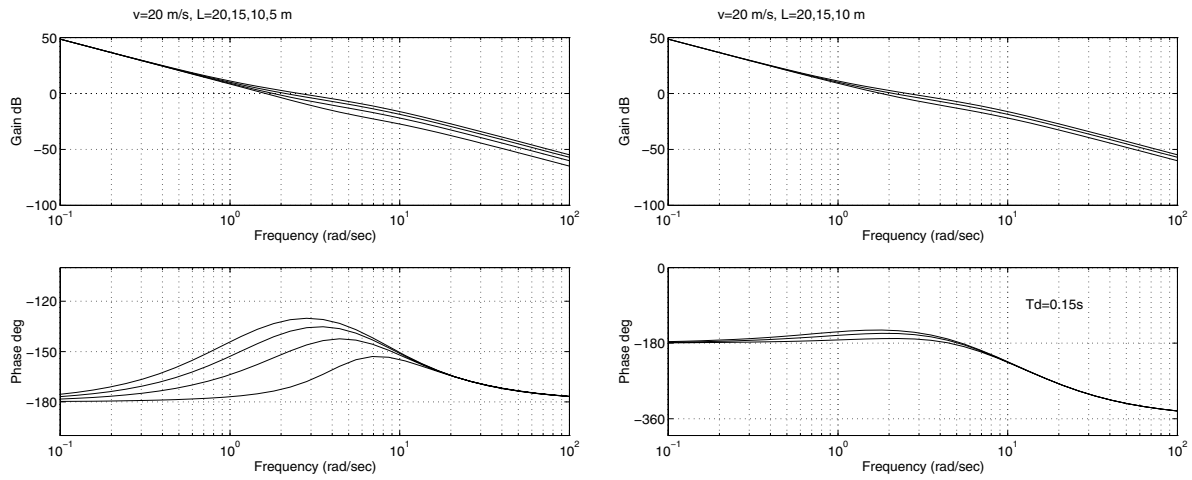

a. Without Delay b. With Delay

Figure 4: (a) Bode plot $V(s)$ for varying look-ahead L = 5,10,15,20m at v = 20m/s with no delay. Increasing the look-ahead adds substantial phase lead at the crossover frequency. (b) The presence of the delay adds an additional phase lag over the whole range of frequencies. The look-ahead of 20m is still able to provide 27.7° phase margin. When the look-ahead decreases to 10m the phase margin in the presence of delay diminishes and the system becomes unstable. Choosing larger look-ahead is more crucial in the presence of delay.

2 Lane Recognition

The lane recognition module is responsible for recovering estimates for the position and orientation of the car within the lane from the image data acquired by a forward looking CCD video camera. This camera is mounted inside the passenger compartment near the rear view mirror as shown in Figure 5 . The roadway is modeled as a flat surface which implies that there will be a simple projective relationship between the coordinates of points on the image plane and the coordinates of their correspondents on the ground plane [23]. This relationship is captured in Equation (10) where the image plane coordinates are denoted by (u, v) and the ground plane coordinates are denoted by (x, y).

$$\begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \propto H \begin{pmatrix} u \\ v \\ 1 \end{pmatrix} \tag{10}$$

The 3 by 3 homography matrix, H, can be recovered through an offline calibration procedure. This model is adequate for our imaging configuration where a camera with a fairly wide field of view (approximately 30 degrees) is used to monitor the area immediately in front of the vehicle (4 - 25

Figure 5: View from inside the automated Honda accord showing the mounting of the cameras

meters).

The first stage of the lane recognition process is responsible for detecting and localizing possible lane markers on each row of the input image. The lane markers are modeled as white bars of a particular width against a darker background. Regions in the image which satisfy this intensity profile can be identified through a template matching procedure. It is important to remember that the width of the lane markers in the image changes linearly as a function of the image row. This means that different templates must be used for different rows.

Once a set of candidate lane markers has been extracted, a robust fitting procedure based on the Hough transform is used to find the best fit straight line through these points on the image plane. A robust fitting strategy is essential in this application because on real highway traffic scenes the feature extraction procedure will almost always return extraneous features that are not part of the lane structure. These extra features can come from a variety of sources, other vehicles on the highway, shadows or cracks on the roadway, other road markings etc. and can confuse naive estimation procedures based on least squares techniques.

The Hough transform procedure considers a set of candidate straight lines on the image plane and computes a score for each one which indicates how well the line conforms to the lane markers. The contribution of a given image measurement to this score is based upon the distance between the edge marker and the candidate line. The candidate line with the best overall score is returned by the lane recognition system. ¿From these measurements it is a simple matter to compute an

a. 

b. 

Figure 6: These figures show the performance of the lane extraction system on a typical input image

estimate for the lateral position and orientation of the vehicle with respect to the roadway at a particular lookahead distance, L, by making use of Equation (10).

The lane finding system is implemented on an array of TMS320C40 digital signal processors which are hosted on the bus of an Intel-based industrial computer. The system processes images from the video camera at a rate of 30 frames per second with a latency of 57 milliseconds. This latency refers to the interval between the instant when the shutter of the camera closes and the instant when a new estimate for the vehicle position computed from that image is available to the control system. This system has been used quite successfully in all of our experiments and was particularly adept at finding difficult lane markings like "Bott's dot" reflectors on a concrete surface (see Figure 6).

3 Observer Design

In order to estimate the curvature of the roadway we have chosen to implement an observer based on a slightly simplified version of the dynamic system given in Equations 5 and 6. More specifically, in this formulation the vehicles lateral velocity, v_y, is neglected and the yaw rate $\dot{\psi}$ is treated as an input. The resulting system is given in Equation (12).

$$\dot{\boldsymbol{x}}' = A'(v_x)\boldsymbol{x}' + B'\dot{\psi} \tag{11}$$

$$\boldsymbol{y} = C'\boldsymbol{x}' \tag{12}$$

where $\boldsymbol{x}' = [y_L, \varepsilon_L, K_L]^T$, $\boldsymbol{y}' = [y_L, \varepsilon_L]^T$, $A'(v_x) = \begin{bmatrix} v_x & -L & 0 \\ 0 & 0 & v_x \\ 0 & 0 & 0 \end{bmatrix}$, $B' = \begin{bmatrix} -L \\ -1 \\ 0 \end{bmatrix}$ and $C' = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix}$.

Note that the state vector $\boldsymbol{x}'$ includes the road curvature K_L. This differential equation can be converted to discrete time in the usual manner by assuming that the yaw rate, $\dot{\psi}$, is constant over the sampling interval T.

$$\boldsymbol{x}'(k+1) = \Phi(v_x)\boldsymbol{x}'(k) + \beta\dot{\psi} \tag{13}$$

Equation (13) allows us to predict how the state of the system will evolve between sampling intervals.

Measurements are obtained from two sources: the vision system provides us with measurements of y_L and ε_L, while the on-board fiber optic gyro gives us measurements of the yaw rate of the vehicle, $\dot{\psi}$. Our use of the yaw rate sensor measurements is analogous to the way in which information from the proprioceptive system is used in animate vision. The measurement vector $\boldsymbol{y}'$ is used to update an estimate for the state of the system $\hat{\boldsymbol{x}}'$ as shown in the following equation:

$$\hat{\boldsymbol{x}}'^{+}(k) = \hat{\boldsymbol{x}}'^{-}(k) + L(\boldsymbol{y}'(k) - C\hat{\boldsymbol{x}}'^{-}(k)) \tag{14}$$

where $\hat{\boldsymbol{x}}'^{-}(k)$ and $\hat{\boldsymbol{x}}'^{+}(k)$ denote the state estimate before and after the sensor update respectively.

The gain matrix L can be chosen in a number of ways [28], depending on the assumptions one makes about the availability of noise statistics and the criterion one chooses to optimize. In our case, the gain matrix was chosen to minimize the expected error of our estimate in the steady state using the function `dlqe` available in Matlab. The covariances of both the process and measurement noise were estimated by analyzing the data collected by our sensors during trial runs with the vehicle.

4 Controllers

The goal of all of the control schemes presented in the sequel is to track the roadway by regulating the offset at the lookahead, y_L, to zero. Passenger comfort is another important design criterion and this is typically expressed in terms of jerk corresponding to the rate of change of acceleration. For a comfortable ride, no frequency above 0.1-0.5 Hz should be amplified in the path to lateral acceleration [29]. Additional performance criteria may be specified in terms of the maximal allowable offset y_{Lmax} as a response to a step change in curvature and in terms of bandwidth requirements on the transfer function $F(s) = \frac{y_L(s)}{K_L(s)}$ between the offset at the lookahead and the road curvature.

4.1 Lead-lag Control

Analysis of the transfer function given in Equation (9) revealed that at speeds of up to 15 m/s with a lookahead of around 10 meters one can guarantee satisfactory damping of the closed loop poles of $V(\boldsymbol{s})$ and compensate for the processing delay of the vision system using simple unity feedback control with proportional gain in the forward loop. As the velocity increases, the poles of the transfer function move toward the real axis and become more poorly damped which introduces additional phase lag in the frequency range 0.1-2 Hz. Since further increasing the lookahead does not improve the damping, gain compensation alone cannot achieve satisfactory performance. A natural choice for obtaining additional phase lead in the frequency range 0.1-2 Hz would be to introduce some derivative action, however, in order to keep the bandwidth low an additional lag term is necessary. One satisfactory lead-lag controller has the following form:

$$C(\boldsymbol{s}) = \frac{0.09\boldsymbol{s} + 0.18}{0.025\boldsymbol{s}^2 + 1.5\boldsymbol{s} + 20} \tag{15}$$

where $C(\boldsymbol{s})$ is a lead network in series with a single pole. The above controller was designed for a velocity of 30 m/s (108 km/h, 65 mph), a lookahead of 15 m and 60 ms delay. The resulting closed loop system has a bandwidth of 0.45 Hz with a phase lead of 45° at the crossover frequency. A discretized version of the above controller taking into account the 33 ms sampling time of the vision system was used in our experiments.

Since increasing the speed has a destabilizing effect on the vehicle transfer function, $V(\boldsymbol{s})$, designing the controller for the highest intended speed guarantees stability at lower speeds and achieves satisfactory ride quality. In order to tighten the tracking performance at lower speeds individual controllers can be designed for various speed ranges and gain scheduling techniques used to interpolate between them.

4.2 Full State Feedback

Given that the vehicle can be modeled as a linear dynamical system it seems natural to consider standard full state linear feedback laws of the form $\boldsymbol{u} = K\boldsymbol{x}$. A controller was designed for velocity of 20 m/s and a lookahead of 6 meters. The gain matrix, K, was chosen using pole placement techniques such that the two poles of the system that were originally at the origin were moved to a conjugate pair with a damping ratio $\xi = 0.707$ and a natural frequency $\omega_n = 0.989$ rad/s. The other two poles of the system were left unchanged. These pole locations were chosen so that the resulting system would satisfy our step response and bandwidth requirements. Since it is assumed that the full state of the system can be estimated at each instant, a smaller lookahead distance can be employed in this design without sacrificing stability.

In the resulting linear control law, the gain associated with the lateral velocity term v_y was small so we chose to neglect this component of the controller. Estimates for the remaining state variables, y_L, ε_L, and $\dot{\psi}$ are obtained from our observer and the yaw rate sensor.

One problem with this controller design is that it fails to account for the latency of the vision system. These types of delay elements are difficult to account for in a state space formulation. One way to compensate for the latency is through the use of a Smith Predictor which would use the delayed estimate for the system state and the system model to estimate the current state of the system. Unfortunately, this approach is notoriously sensitive to errors in the model.

4.3 Input-Output Linearization

Input-ouput linearization is typically used to linearize nonlinear systems by state feedback as described in [35]. The application of this technique to our bicycle model is not, strictly speaking, linearization by state feedback since the bicycle model is already linear. Nonetheless, this technique can be applied to render the model independent of the vehicles longitudinal velocity, v_x. In this case the feedback law has a zero canceling effect instead of a linearizing one and makes the vehicle dynamics poles unobservable.

If the bicycle model of Equation 7 is rewritten in the form:

$$\dot{\boldsymbol{x}} = f(\boldsymbol{x}) + g(\boldsymbol{x})\delta_f \tag{16}$$

$$y_L = h(\boldsymbol{x}) \tag{17}$$

the control law required to linearize this system can be obtained by differentiating the y_L output twice with respect to time. [1] The resulting control law has the form given in equation 19

$$\delta_f = \frac{1}{L_g L_f^1 h(x)}(-L_f^2 h(x) + u) \tag{18}$$

$$= \frac{-1}{b_1 + lb_2}((\frac{l}{I_z v_x}(l_r c_r - l_f c_f) - \frac{1}{m v_x}(c_f + c_r))v_y + (\frac{1}{m v_x}(c_r l_r - c_f l_f) - \frac{l}{I_z v_x}(l_f^2 c_f + l_r^2 c_r))\dot{\psi} + u) \tag{19}$$

where L_g^i denotes the i-th Lie derivative along g.

Employing this control law yields a second order equation of the form $\ddot{y_L} = u$. Once the system has been reduced to this form we can employ the same lead-lag control law described in Section 6.1 to compute a control input u which will stabilize the system and achieve the desired performance goals.

[1] Two differentiations are required since the relative degree of the system is 2

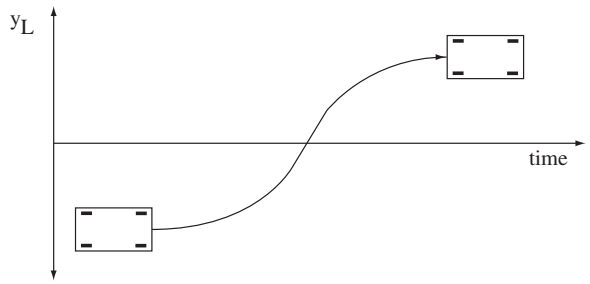

Figure 7: Lane change maneuver

4.4 Feedforward Control

The estimate for road curvature returned by the observer can be used as part of a feedforward control strategy. The steady state steering input, δ_{ref}, that is required to track a reference curvature, K_{Lref}, can be computed from the state equations by setting $[\dot{v}_y, \ddot{\psi}, \dot{y}_L, \dot{\varepsilon}_L]^T$ to 0.

$$\delta_{ref} = K_{ref}\,(l - \frac{(l_f c_f - l_r c_r) v_x^2 m}{c_r c_f l})\ . \tag{20}$$

This feedforward control component can be added to any of the control schemes that have been described. The feedforward control law allows the system to anticipate changes in curvature ahead of the car and improves the transient behavior of the vehicle when entering and exiting curves. The effectiveness of the feedforward term will, of course, depend on the quality of the curvature estimates supplied by the observer.

4.5 Lane Change Maneuvers

Lane change maneuvers are accomplished by supplying a reference trajectory, $y_L(t)$, as an input to the lateral control systems. This reference trajectory is a simple fifth order spline which smoothly moves the vehicle from one lane to another as shown in in Figure 7. The curvature of the reference trajectory is also supplied as an additive input to the feedforward control law.

5 Vision for longitudinal vehicle control

Another modality of the automated vehicle is a longitudinal control system which combines both laser radar and vision sensors, enabling throttle and brake control to maintain a fixed distance to a lead car. Currently, steering control can be driven by a combination of lane tracking described in the previous section and magnetometers detecting magnetic "nails" in the road, and longitudinal control achieved using the laser radar sensor. In our work we explored possibility of using vision

which can potentially provide higher bandwidth (30/60Hz) output than is available from the laser radar system [9].

Our approach combines shape reconstruction (2D planar reconstruction rather than the usual 3D) from stereo/motion with motion estimation, using recently developed robust and efficient feature matching methods. The resulting algorithm runs at 3-5Hz, and frame-rate performance (30Hz) is achieved by combining this with a normalized cross-correlation algorithm for stereo and motion computation, which runs in parallel. Employing these quite sophisticated tools we are capable to effectively track a vehicle over an extended time, i.e. a period measured in minutes rather than frames.

We shall assume that the viewed vehicle is an affine projection of a planar object. We shall not assume any knowledge of the camera calibration for the purpose of tracking the vehicle. Thus we have possibly the simplest available imaging model. The lead vehicle presents its rear end to the cameras, with little change in orientation, which might otherwise induce the 3D cues that we are ignoring. The depth relief is also small, since only the rear of the car is visible. The smallest range we are considering in our experiments is 10m, so that cars of size 2m will subtend an angle of at most 10^{o}, justifying the assumption of parallel projection from scene to image. The results at the end of this paper will demonstrate that these minimalist assumptions are appropriate.

5.1 Visual tracking

In order to track an object over an extended time, it is necessary to compute the *position* of the object, which in this context may only be computed by integrating the velocity over time. Since there are inevitable errors in the computed velocities, these errors will tend to accumulate over time. Thus we can expect the computed position to drift. The drift problem can be eliminated in correlation trackers by fixing the template used for correlation, converting it into a position-based tracker; however then the tracker will work only for a short time, because relatively small motions, especially rotations, will break the tracker, although Shi & Tomasi [18] suggest a partial solution to that problem through an affinely deformable template. Correlation with a fixed template forms the lower level of our tracker, and we use it in such a way that a single template is used for only a short period of time.

Feature tracking algorithms have the capability to allow the position of an object to be accurately estimated over an extended time. This aspect is of vital importance in the context of sensing for control, where the sensor is required to return accurate error feedback during the whole period of the control task. Moreover we have demonstrated in previous work [1] that vehicle tracking using features can be made robust both to partial occlusion of the vehicles and to lighting changes in

the environment. The major problems that have to be overcome when tracking features are the fragmentary nature of the data (features appear, disappear and change shape) and the integration of feature data from multiple images in a statistically valid manner. We now describe briefly some of the problems. More detailed description can be found in [13].

5.2 Fixation/Scene Reconstruction

In [17] a fixation technique was described that allows a single "fixation point" to be chosen from a cluster of tracked features, in a way that is robust to losing track of individual features, while allowing the same *object* point to be fixated over time.

While we have a simpler problem, with a fixed camera and a 2D scene representation, similar principles apply. The "fixation point" here refers to the point chosen as representative of the vehicle for the purpose of estimating its position, and hence its range. Maintaining scene structure explicitly within a reconstruction algorithm stabilizes the computation of motion over time, so a logical extension of the fixation point transfer algorithm is to recursively update the structure of the tracked object, and employ the improved motion estimates to perform fixation point transfer. This is the method we have implemented. The reconstruction technique detailed below is a 2D affine version of the Variable State Dimension Filter (VSDF) algorithm [14]. The VSDF is a general algorithm for visual reconstruction that deals naturally with fragmentary data and combines data from multiple images in a near-optimal manner.

5.3 2D Affine Reconstruction

The 2D affine projection from scene to image can then be written in its most general form as

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} M_{1X} & M_{1Y} \\ M_{2X} & M_{2Y} \end{pmatrix} \begin{pmatrix} X \\ Y \end{pmatrix} + \begin{pmatrix} t_x \\ t_y \end{pmatrix} \quad \text{or} \quad \mathbf{z} = M\mathbf{X} + \mathbf{t}. \tag{21}$$

Here $\mathbf{z}$ is the image point, $\mathbf{X}$ is a 2D scene point, and M, $\mathbf{t}$ constitute the camera matrix parameters, which we will term the *motion* parameters because they represent the camera/scene motion over time. The 2D affine reconstruction problem for point features is: given $\mathbf{z}_i^{(j)}$ for multiple features i in multiple images j, determine the motion $M^{(j)}$, $\mathbf{t}^{(j)}$ and structure $\mathbf{X}_i$. We employ the variable state dimension filter algorithm which achieves virtually the same accuracy as previously used batch algorithms [20], but has the advantages of being recursive, not requiring complete data, and allowing new features to be added to the reconstruction as they appear and discarded features to be removed.

Given the recursively computed scene reconstruction, the transfer of the fixation point into the new stereo pair with computed motion parameters $M_l, \mathbf{t}_l$ and $M_r, \mathbf{t}_r$ is simply $\mathbf{z}_{fl} = M_l \mathbf{X}_f + \mathbf{t}_l, \quad \mathbf{z}_{fr} = M_r \mathbf{X}_f + \mathbf{t}_r$. The transferred positions $\mathbf{z}_{fl}$, $\mathbf{z}_{fr}$ are now converted via triangulation to 3D world coordinates to be interpreted as a range measurement. We choose $\mathbf{X}_f$ to be the centroid of the structure vectors computed from the initial batch computation of the 2D reconstruction.

5.4 Stereo/Temporal Matching

A major issue with all reconstruction techniques is their reliance on high quality, essentially outlier-free input data. The method that has recently been proposed to achieve this is to select a set of feature matches that are globally consistent, in the sense of satisfying the rigidity constraint.

We follow [21] and apply the RANSAC algorithm of Fischler & Bolles [7] to compute a large subset of feature matches consistent with a single set of 2D transformation parameters. Part of the stereo matching algorithm is also an algorithm for enforcing uniqueness of individual matches. Both algorithms are described in more detail in [13].

5.5 Layered Tracking

The frame-rate (30Hz) performance of the tracker results from the coordination of two separate tracking algorithms, a frame-rate correlator which computes the motion in both images, and the corner feature-based fixation algorithm, which runs at 3-5Hz depending on the size of the region used for the corner detection (maximum 140×100 pixels). The two processes run on separate C40s, and are coordinated in such a way that the correlator is always using an image region centered around the latest fixation point for the correlation. The two processes communicate whenever the fixation algorithm has any results to pass on. In addition, the laser radar provides the fixation algorithm with the initial bounding boxes around the vehicle, so the radar may be considered a third layer of the tracker. This provides the system with a great deal of robustness. If the correlator fails for any reason (usually due to not finding a motion with a high enough correlation score) it simply waits for the fixation algorithm to provide it with a new template/position pair. If the fixation algorithm fails, then it also must wait for the laser to provide it with a new bounding box pair.

6 Implementation

Figure 8 shows the major components of our autonomous vehicle control system which was implemented on the Honda Accord LX shown in Figure 9. This system takes input from a range

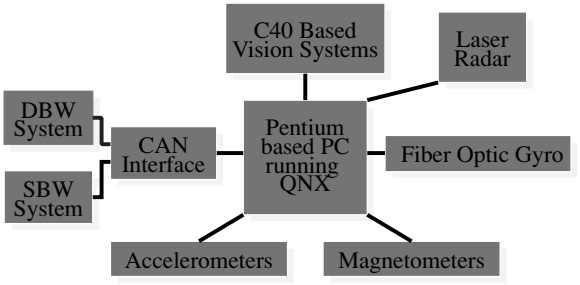

Figure 8: System Diagram

Figure 9: The Honda Accord LX sedan used in our experiments

of sensors which give it information about its own motion, (speedometer, yaw rate sensor and accelerometers), its position in the lane, (vision system and magnet nail sensors), and its position with respect to other vehicles in the roadway (the laser radar system).

All of these sensor systems were interfaced to an Intel-based industrial computer which ran the QNX real time operating system. All of the control algorithms and most of the sensor processing were performed by the host computer. The real-time lane extraction operation was carried out on a network of TMS320C40 Digital Signal Processors which was hosted on the bus of the main computer.

The experimental setup for the vision-based tracker for longitudinal control is illustrated in Figure 10. The off-line version of the tracking algorithm was tested on approximately 20 minutes of synchronized video and laser radar [9].

7 Experimental Results

In order to compare the various feedback strategies we implemented them on our experimental vehicle and collected data from a number of trial runs. Our test track was a 7 mile oval (see Figure 11) and our experiments were run at speeds of approximately 75mph to simulate actual

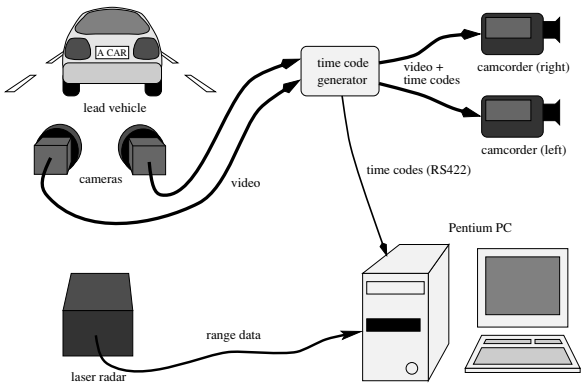

Figure 10: The experimental setup.

highway conditions. Each experimental trial lasted at least 5 minutes, long enough to explore how each controller fared on the straight sections, the curved sections and the transitions between them. Figures 11 and 12 describe the performance of tested control strategies without and with the feedforward term.

Figures 11a, 11b and 11c indicate the tracking performance of the lead-lag, full state feedback and I/O linearization controllers respectively, that is they indicate the offset of the centerline of the road at a distance of 15 meters ahead of the vehicle in case of lead-lag and I/O linearization and 6 meters in case of full state feedback controller. Since the controllers are designed to regulate this quantity to zero, this is an appropriate value to monitor.

Figures 11d, 11e and 11f indicate the velocity profiles during these runs while Figures 11g, 11h and 11i denote the lateral acceleration experienced at the center of gravity of the vehicle. The plots indicate a steady state offset for all of the controllers in the curved sections of the track; this is expected since all of the controllers have to produce a non-zero steering control effort on these sections based on feedback. The lead-lag controllers tracking performance is superior to that of the other two control strategies. For the full-state feedback controller there is a noticeable overshoot during transition between curved and straight segments and its performance degrades when the velocity increases above the value considered in the design. One possible approach to improving the transient behavior of this controller would be to increase the lookahead distance used in the design. Because the lookahead distance used for the full state feedback controller was smaller the offset measurements are less noisy. The tracking performance of the I/0 linearized controller is quite good at lower velocities but the ride becomes a little rougher at higher velocities.

The plots in Figure 12 demonstrate the effect of the feedforward control term on the overall tracking performance for all tested control strategies. The first row of plots indicates the tracking performance measured in terms of the offset at the lookahead, the second row depicts the curvature

estimate used in the feedforward term, which was provided by the observer, the third row denotes the velocity profiles of the experiments and the last row shows the lateral acceleration profiles. Notice that the steady state offset in the curved sections was essentially eliminated. The offset plots all exhibit a slight overshoot during transitions in curvature until the curvature estimates converge. The lateral acceleration profile of the Input/Output linearizing controller is somewhat better than that of the other two indicating a smoother ride in this case. In the case of full state feedback controller the spikes in the offset measurements and the lateral acceleration profile correspond to the lane change maneuvers which the vehicle performed at lower speeds (50 mph).

For the tracking algorithm we selected an initial window surrounding the lead vehicle, although subsequent processing was completely automatic. Figure 13 show some example images, with the tracking results superimposed. The corner features are shown as small crosses, white for those matched over time or in stereo, and black for unmatched features. The black and white circle indicates the position of the fixation point, which ideally should remain at the same point on the lead car throughout the sequence. The white rectangle describes the latest estimate of the bounding box for the vehicle.

We have attempted here to summarize significant aspects of our data. Images 1 and 2 show the first stereo pairs in the sequence, where the vehicle is close (17m) to the camera and range estimates from stereo disparity may be expected to be accurate. By contrast images 421 and 422 were taken when the vehicle was 60m from the camera (the greatest distance achieved during the sequence). Here we can predict that depth estimates from stereo will be unreliable, since the disparity relative to infinity is only a few pixels and so difficult to measure. It will still be feasible to use the change in apparent size measured by the motion processing to obtain reasonable range estimates. We computed the range and bearing estimated from the laser radar range finder and plot them together with the corresponding data collected from the vision algorithms in figures 14 and 15. Depth from stereo is computed by inverting the projection of the fixation point at each image pair and finding the closest point of intersection of the two resulting space rays. The cameras are very roughly calibrated.

The results for the real-time version of the algorithm, in a simple scenario wherein the lead car is stationary and the following car moves backwards and forwards, are shown in figure 16. Depth from stereo is computed by inverting the projection of the fixation point at each image pair and finding the closest point of intersection of the two resulting space rays. Figure 16 shows the comparison of the vision and laser radar data in a 200 second segment of data. There is clearly an offset between the graphs, due to a calibration error, which we could correct fairly easily. The vision data is quite

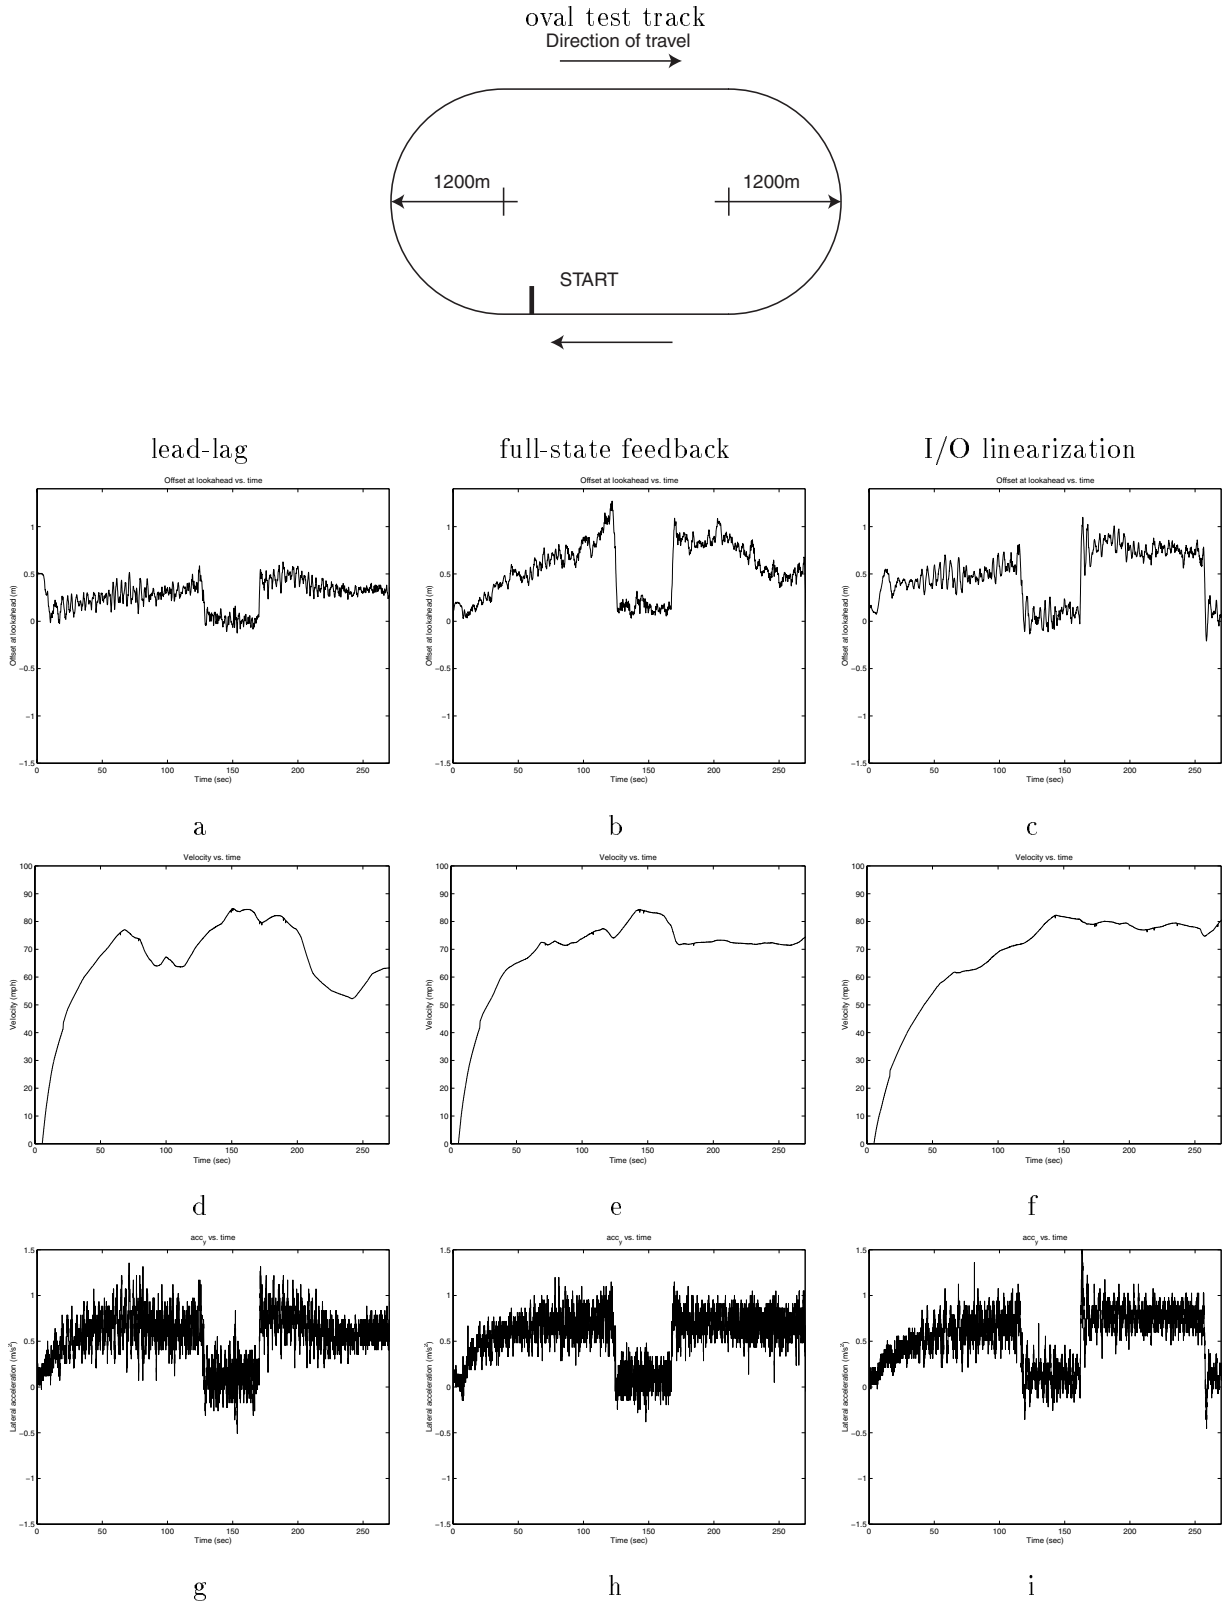

Figure 11: This figure presents a side by side comparison of the results obtained when our test vehicle was driven on an oval track under each of the control schemes that were implemented. The first column of plots correspond to data collected under the lead-lag control scheme, the second to the full-state feedback controller and the third to the I/O linearization method. Figures a, b and c indicate the tracking performance of the controllers, that is they indicate the offset of the centerline of the road at a distance of 15 meters ahead of the vehicle in the case of lead-lag and I/O

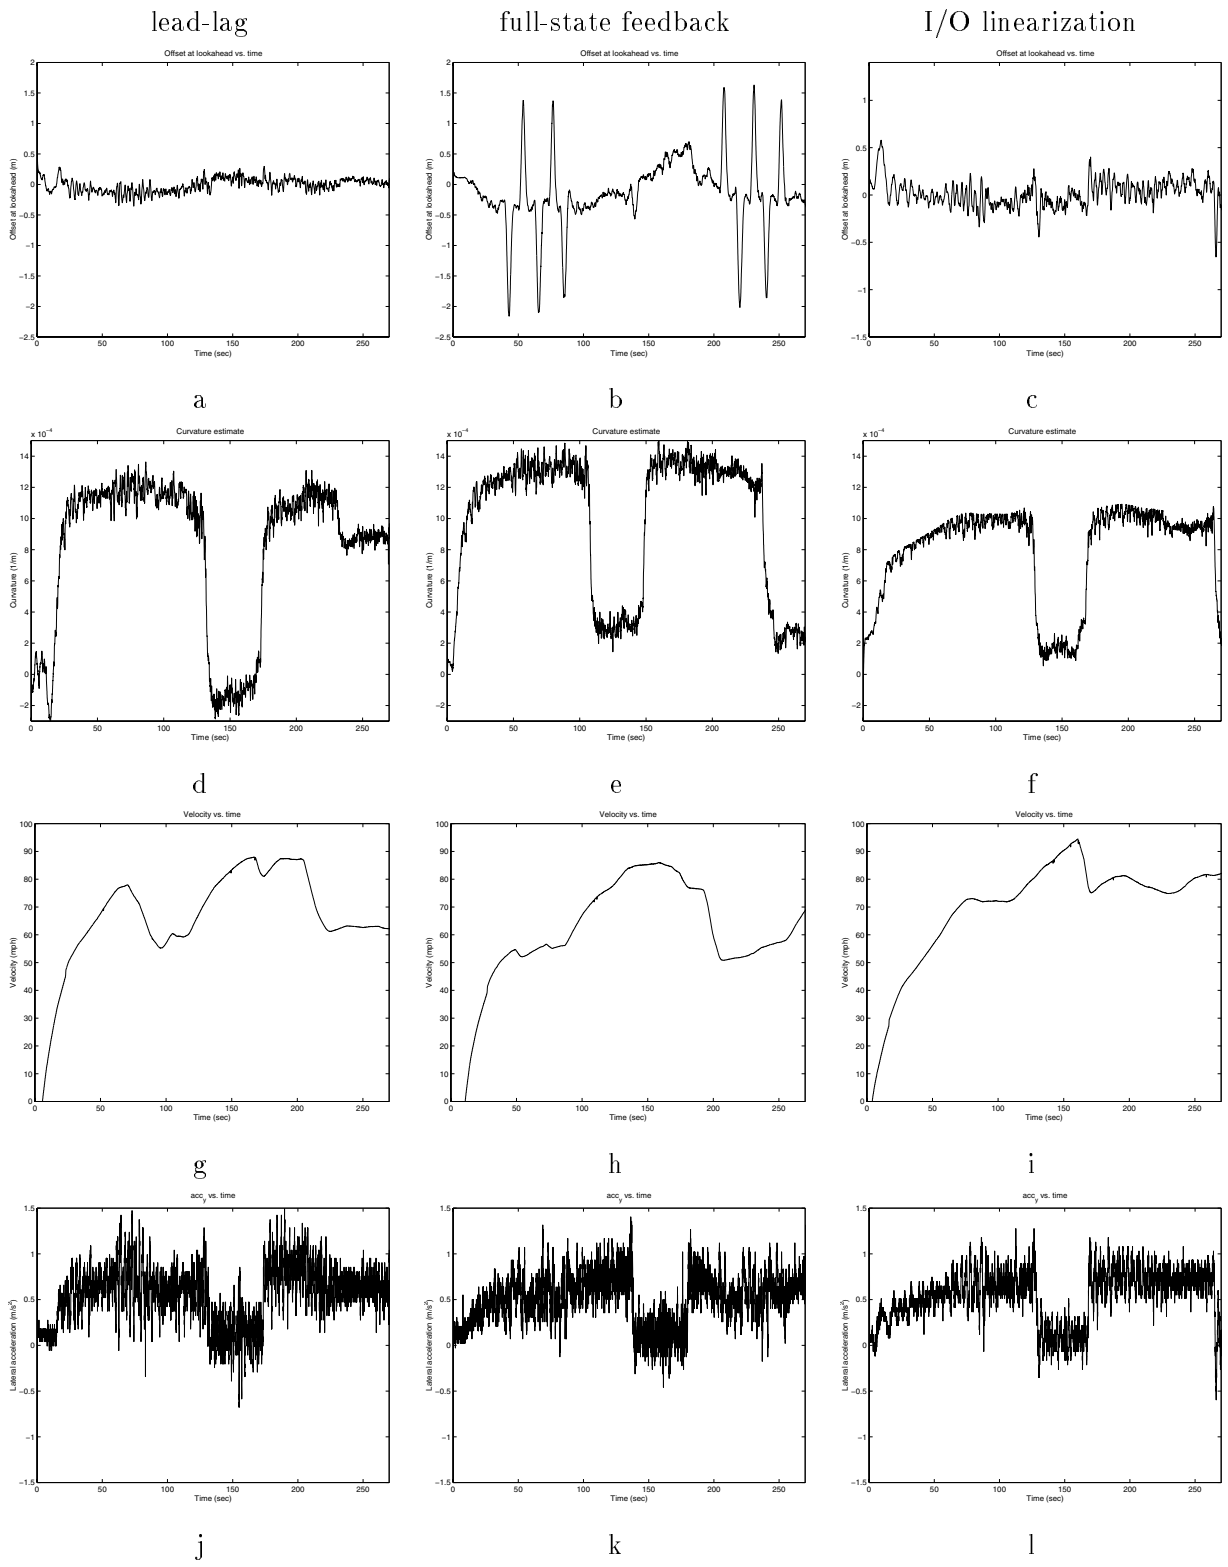

Figure 12: These plots demonstrate the effect of the feedforward control term on the overall tracking performance for all tested control strategies. The first row of plots indicates the tracking performance measured in terms of the offset at the lookahead, the second row depicts the curvature estimate used in the feedforward term, which was provided by the observer, the third row denotes the velocity profiles of the experiments and the last row shows the lateral acceleration profiles. Notice that the steady state offset in the curved sections was essentially eliminated. The offset

noisy; this is partly due to the small baseline (26cm) used for triangulation to obtain the range measurements, and again this could be improved by increasing the baseline. A comparison of the heading angle is shown in figure 17.

The software was written using the C library "Horatio" which has been developed at Oxford and Berkeley for the purpose of supporting efficient computer vision applications. HTML documentation for Horatio may viewed from the first author's WWW home page at `http://www.cs.berkeley.edu/~pm/`, and the complete library may also be downloaded from that site.

8 Conclusions

This report has presented an analysis of the vision-based lateral control task and an investigation of how the characteristics of this problem change as a function of important system parameters such as vehicle velocity, lookahead distance and processing delay. We have also discussed the results of our experiments with three different feedback control strategies; lead-lag control, full state feedback and input-output linearization. Our experiments indicate that all three of the feedback control strategies that we implemented provided acceptable performance on the lateral control task with the lead lag control law yielding the best tracking performance of the three. The data also shows that the curvature feedforward component definitely improves the tracking performance of all three control strategies. It allows the system to eliminate steady state tracking errors when following a curve and it minimizes the transient response of the system to changes in curvature.

The strategy behind the design of the feedback control laws was based on the observation that the behavior of our system was dominated by the two poles at the origin, since the other two poles are well behaved as long as the lookahead distance is large enough. This allowed us to design controllers for the highest intended operating velocity, which would operate satisfactorily in the whole range of lower velocities. However this approach sacrifices some performance criteria at lower velocities.

We have demonstrated an approach for vision based lateral and longitudinal vehicle control. The vision based lane tracking system used a robust fitting strategy which allowed us to overcome

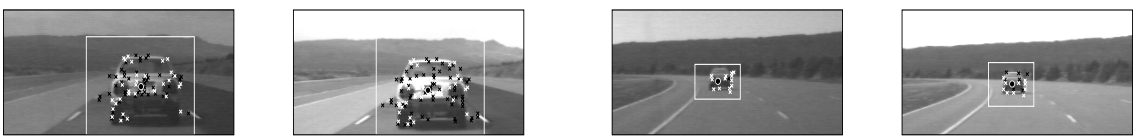

Figure 13: Example stereo-pairs from the tracking sequence.

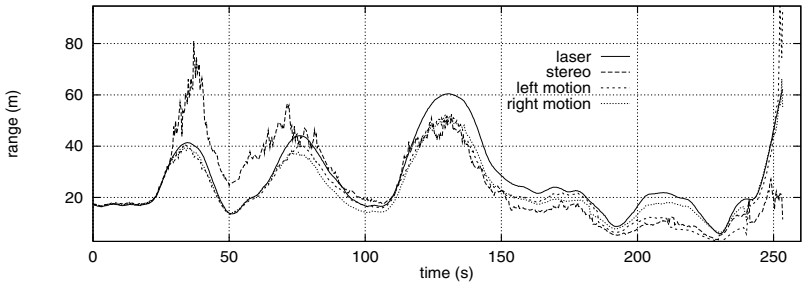

Figure 14: Comparison of range estimates from laser radar and vision.

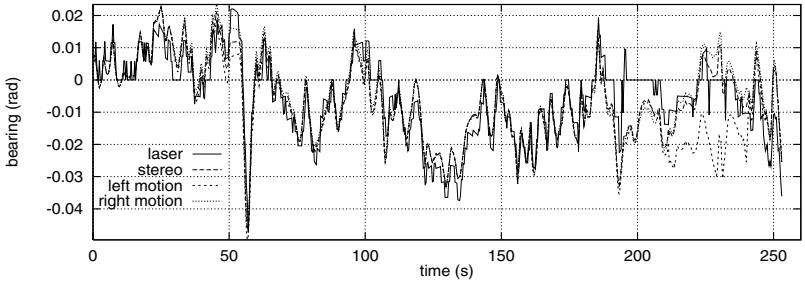

Figure 15: Comparison of bearing estimates from laser radar and vision.

spurious lane markings and other distracting features that are common on CALIFORNIA highways. Closed-loop simulations and careful analysis revealed the importance of the look-ahead information provided by the vision system, both in the presence of the delay caused by processing of the visual information and changes in the vehicle dynamics with increasing speed. The delay plays an important role in the system and should be taken into account explicitly in case of output feedback strategies, such as the ones we presented. We showed that sufficiently large look-ahead and appropriate choice of gain can compensate for the additional phase lag introduced by delay and vehicle dynamics at lower velocities. At higher velocities additional lead action was introduced in order to achieve desired phase margin. Introducing a real-time observer process into the system not only reduced the noise inherent in the system's sensor measurements, but also provided an accurate estimate of the current vehicle state, circumventing the delay in the vision system and permitting the implementation of more advanced state-space based controllers. The curvature of the road was incorporated into the observer process and the estimates were used for feedforward control strategies. The presence of the feedforward term improved tracking performance in curved road segments.

The information about the vehicle ahead for longitudinal control was provided by a stereo vision algorithm in conjunction with a scanning laser radar sensor. The vision algorithm is built on

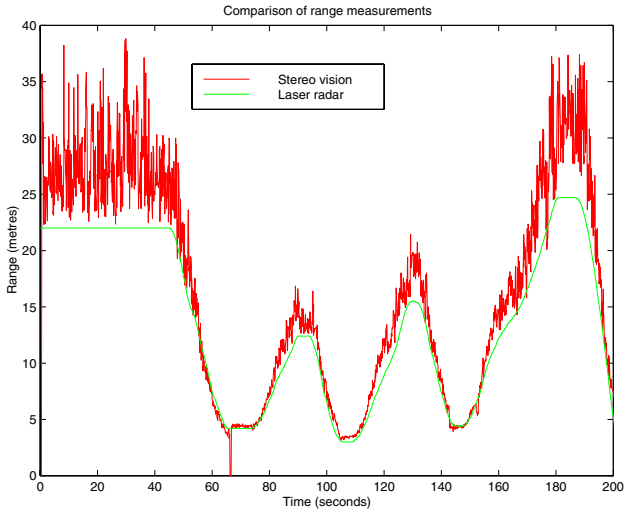

Figure 16: Comparison of range estimates between stereo vision and laser radar from a real-time run.

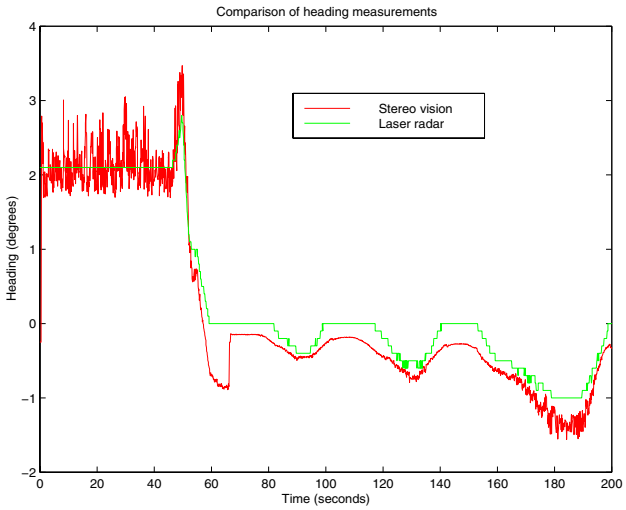

Figure 17: Comparison of heading estimates between stereo vision and laser radar from a real-time run..

fixation and reconstruction algorithms designed for active vision systems, which combine stereo and motion cues. The current version of the algorithm has been implemented on a network of C40 DSP's and preliminary results indicate that the range estimates are comparable with the laser radar system. The layered approach to vehicle tracking that involves the use of a simple correlation tracker and a more sophisticated temporal/stereo reconstruction algorithm was designed to take advantage of the merits of the two approaches. Correlation tracking is easy to implement in real time, but is not suitable for extended tracking. Temporal/stereo reconstruction is more suited to extended tracking, but it involves a large and variable amount of computation. By utilizing both approaches in parallel, we have achieved a real-time tracker that can maintain tracking and compute the distance to the vehicle for a long period of time.

The resulting system has been tested extensively during the development time and during the NAHSC Demonstration in August 1997 in San Diego when the system was operational for 4 days. We have completed a thorough analysis and comparison of the tested vision based control stategies for lateral control.

References

[1] D. Beymer, P.F. McLauchlan, B. Coifman, and J. Malik. A real-time computer vision system for measuring traffic parameters. In *Proc. of the IEEE Conf. on Computer Vision and Pattern*

Recognition, 1997.

[2] A. Blake, R. Curwen, and A. Zisserman. A framework for spatiotemporal control in the tracking of visual contours. *International Journal of Computer Vision*, 11(2):127–146, October 1993.

[3] P.J. Burt, J.R. Bergen, R. Hingorani, R. Kolczynski, W. A. Lee, A. Leung, J. Lubin, and H. Shvaytser. Object tracking with a moving camera. In *Proc. IEEE Workshop on Visual Motion*, pages 2–12, 1989.

[4] T.H. Cormen, C.E. Leiserson, and R.L. Rives. *Introduction to Algorithms*. MIT Press, 1994.

[5] E.D. Dickmanns and B.D. Mysliwetz. Recursive 3-d road and relative ego-state recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):199–213, February 1992.

[6] S.M. Fairley, I.D. Reid, and D.W. Murray. Transfer of fixation for an active stereo platform via affine structure recovery. In *Proc. 5th Int'l Conf. on Computer Vision, Boston*, pages 1100–1105. IEEE Computer Society Press, 1995.

[7] M.A. Fischler and R.C. Bolles. Random sample concensus: A paradigm for model fitting with applications to image analysis and automated cartography. *Comm. ACM*, 24(6):381–395, 1981.

[8] C. Harris. Tracking with rigid models. In A. Blake and A. Yuille, editors, *Active Vision*. MIT Press, Cambridge, MA, 1992.

[9] H. Kikuchi, M. Ishiyama, and T Nakajima. Development of laser radar for radar brake system. In *Proceedings of the International Symposium on Advanced Vehicle Control*, pages 385–389, Tsukuba Research Centre, AIST, MITI, Japan, 1994.

[10] D. Koller, J. Weber, and J. Malik. Robust multiple car tracking with occlusion reasoning. In *Proc. 3rd European Conf. on Computer Vision, Stockholm*, volume 1, pages 189–196, May 1994.

[11] J. Košecká, R. Blasi, C.J. Taylor, and J. Malik. Vision-based lateral control of vehicles. In *Proc. Intelligent Transportation Systems Conference, Boston*, 1997.

[12] P. F. McLauchlan and J. Malik. Vision for longitudinal control. In A. Clark, editor, *Proc. 8th British Machine Vision Conf., Essex*. BMVA Press, 1997.

[13] P. F. McLauchlan and J. Malik. Vision for longitudinal vehicle control. In *Proc. Intelligent Transportation Systems Conference, Boston*, 1997.

[14] P.F. McLauchlan and D.W. Murray. A unifying framework for structure and motion recovery from image sequences. In *Proc. 5th Int'l Conf. on Computer Vision, Boston*, June 1995.

[15] P.A.Beardsley, A.Zisserman, and D.W.Murray. Sequential updating of projective and affine structure from motion. *International Journal of Computer Vision*, 23(3), 1997.

[16] K. Pahlavan, T. Uhlin, and J.O. Eklundh. Dynamic fixation and active perception. *IJCV*, 17(2):113–135, February 1996.

[17] I. D. Reid and D. W. Murray. Active tracking of foveated feature clusters using affine structure. *International Journal of Computer Vision*, 18(1):1–20, April 1996.

[18] J. Shi and C. Tomasi. Good features to track. In *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition*, pages 593–600, 1994.

[19] S. M. Smith. ASSET-2: Real-Time Motion Segmentation and Shape Tracking. In *Proc. 5th Int'l Conf. on Computer Vision, Boston*, pages 237–244, 1995.

[20] C. Tomasi and T. Kanade. Shape and motion from image streams under orthography: A factorization approach. *International Journal of Computer Vision*, 9(2):137–154, 1992.

[21] P.H.S. Torr, P.A. Beardsley, and D.W. Murray. Robust vision. In E. Hancock, editor, *Proc. 5th British Machine Vision Conf., York*, pages 145–154. BMVA Press, 1994.

[22] P.H.S. Torr, A. Zisserman, and S.J. Maybank. Robust detection of degenerate configurations for the fundamental matrix. In *Proc. 5th Int'l Conf. on Computer Vision, Boston*, pages 1037–1042, 1995.

[23] Z. Zhang and O. Faugeras. *3D Dynamic Scene Analysis*. Springer-Verlag, 1992.

[24] Jitendra Malik, Camillo J. Taylor, Joseph Weber, Dieter Koller and Quang-Tuan Luong. A combined approach to stereopsis and lane finding. UCB-ITS-PRR-97-27, California PATH, final report, 1997.

[25] Robert. S. Blasi. A study of lateral controllers for the stereo drive project. Master's thesis, Department of Computer Science, University of California at Berkeley, 1997.

[26] E. D. Dickmanns and B. D. Mysliwetz. Recursive 3-D road and relative ego-state estimation. *IEEE Transactions on PAMI*, 14(2):199–213, February 1992.

[27] B. Espiau, F. Chaumette, and P. Rives. A new approach to visual servoing in robotics. *IEEE Transactions on Robotics and Automation*, 8(3):313 – 326, June 1992.

[28] Arthur Gelb *et al. Applied optimal estimation.* MIT Press, 1994.

[29] J. Guldner, H.-S. Tan, and S. Patwarddhan. Analysis of automated steering control for highway vehicles with look-down lateral reference systems. *Vehicle System Dynamics (to appear)*, 1996.

[30] Jana Košecká. Vision-based lateral control of vehicles:look-ahead and delay issues. Internal Memo, Department of EECS, University of California Berkeley, 1997.

[31] M. F. Land and D. N. Lee. Where we look when we steer? *Nature*, 369(30), June 1994.

[32] Yi Ma, Jana Košecká, and Shankar Sastry. Vision guided navigation for a nonholonomic mobile robot. In *submitted to CDC'98*, 1997.

[33] Ü. Özgüner, K. A. Ünyelioglu, and C. Hatipoğlu. An analytical study of vehicle steering control. In *Proceedings of the 4th IEEE Conference on Control Applications*, pages 125 – 130, 1995.

[34] H. Peng. *Vehicle Lateral Control for Highway Automation.* PhD thesis, Department of Mechanical Engineering, University of California, Berkeley, 1992.

[35] A. Isidori. *Nonlinear Control Systems.* Springer-Verlag, 1989.

[36] Chuck E. Thorpe, Martial Herbert, Takeo Kanade and Steve Shafer. Vision and navigation for the Carnegie-Mellon Navlab. In *IEEE Trans. Pattern Analysis and Machine Intelligence*, 10(3):342-361, 1988.

[37] D. Raviv and M. Herman. A 'non-reconstruction' approach for road following. In Proceedings of the SPIE, editor, *Intelligent Robots and Computer Vision*, pages 2–12, 1991.

[38] Camillo J. Taylor, Jitendra Malik, and Joseph Weber. A real-time approach to stereopsis and lane-finding. In *Proceedings of the 1996 IEEE Intelligent Vehicles Symposium*, pages 207–213, Seikei University, Tokyo, Japan, September 19-20 1996.

[39] Jitendra Malik, Camillo J. Taylor, Phil McLauchlan, Jana Košecká. Development of Binocular Stereopsis for Vehicle Lateral Control, Longitudinal Control and Obstacle Detection. California PATH-MOU257, final report, 1998.