

A Comparative Study of Vision-Based Lateral Control Strategies for Autonomous Highway Driving

Camillo J. Taylor
Computer and Information Science Dept.
University of Pennsylvania
Philadelphia, PA

Jana Košecká, Robert Blasi, Jitendra Malik
EECS Dept.
U. C. Berkeley
Berkeley, CA

Abstract

With the increasing speeds of modern microprocessors it has become ever more common for computer vision algorithms to find application in real-time control tasks. In this paper we present an analysis of the problem of steering an autonomous vehicle along a highway based on the images obtained from a CCD camera mounted in the vehicle. We explore 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.

1 Introduction

With the increasing speeds of modern microprocessors it has become ever more common for computer vision algorithms to find application in real-time control tasks. In particular, the problem of steering an autonomous vehicle along a highway using the output from one or more video cameras mounted inside the vehicle has been a popular target for researchers around the world and a number of groups have demonstrated impressive results on this control task.

The goal of our research efforts in this field has been to understand the fundamental characteristics of this vision based control problem and to use this knowledge to design better control strategies. In this paper we present an analysis of the problem of vision-based lateral control and describe 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 system.

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. These experiments allowed us to verify the accuracy and efficacy of our modeling and control techniques.

Section 2 of this paper discusses relevant prior work in this area. Section 3 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 4 describes the strategy used to extract lane markings from the video imagery and section 5 describes the design of an observer that we use to estimate the states of our system and the curvature of the roadway. Section 6 describes the various control strategies that we implemented on our experimental platform and section 8 presents the results of the experiments that we carried out with these controllers. Section 9 contains the conclusions that we have drawn from these experiments. A brief description of our experimental vehicle is provided in section 7.

2 Previous Work

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.

There have been several attempts to formulate the vision based steering task in the image plane. Raviv and Herman in [aMH91, HNH⁺97] suggested the use of measurement of the projection of road tangent point and its optical flow in the image for generating steering commands and proposed some simple control strategies. The stability and sensitivity issues of this approach have not been addressed. While the approach worked for many road situations it is not clear whether this cue is sufficient for general road scenarios. The automated steering task using vision has been also formulated within the visual servoing paradigm [ECR92]. A stability analysis was provided, for an omnidirectional mobile base trying to align itself with respect to a straight road. Both of these approaches employ a simple kinematic model of the vehicle and do not address the dynamic behavior of the system which becomes increasingly pertinent at higher speeds (above 20m/s).

Ozguner et al [ÖÜH95, ÖÜH97] investigated the problem of controlling a vehicle with non-trivial dynamics using measurements for the offset of the vehicle with respect to the lane at a certain distance ahead of the vehicle. These authors propose a simple proportional control law and present an analysis which shows that the look-ahead distance can always be chosen large enough to guarantee the closed loop stability of the system, given some limits on longitudinal speed. This analysis holds for a general class of *look-ahead* systems, where the measurements are naturally available ahead of car (e.g. radar, vision). However, they do not consider the destabilizing effects of latencies due to computational delays.

Dickmanns et. al. [DM92] developed a system that drove autonomously on the German Autobahn as early as 1985. This work makes extensive use of the Extended Kalman filter as a framework for integrating measurements from multiple sources over time to produce a coherent estimate of the lateral position and orientation of the vehicle and the curvature of the roadway. The control strategy employed was based on a full state feedback control design where the feedback gains are selected via pole-placement.

The Navlab project at CMU has produced a number of successful visually guided autonomous vehicle systems [THKS88]. The most recent incarnation of this system is based on the Rapid Adaptive Lateral Position Handler (RALPH) [Pom95] which robustly extracts lane markings in the image data obtained from the onboard camera. Bertozzi and Broggi [BB97] describe a system for performing lane extraction on images which makes use of the Inverse Perspective Mapping between the image plane and the road plane.

In the psychophysics literature Land and his colleagues [LL94] studied the correlation between the direction of gaze and the steering performance of human drivers. They asserted that while driving on a curved road the gaze shifts from tangent points on the inside of each curve. The new tangent point is sought for about 1-2s before entering the curve. In addition to previewing the road ahead of the car, authors in [LH95] had shown that drivers also use the information from the near range in front of the car. The information from the near range improves the position of the car in the lane. At higher speeds the preview information becomes more important. The preview information has been shown sufficient for satisfactory road following in spite of the presence of delay ≈ 0.73 s between gaze-shifts and steering movements. Part of the delay was attributed to time it takes to process visual information [Lan96]. In addition to studying sensing aspects of human drivers, several models of driver steering have been developed [Hes90]. The understanding of human steering performance gained from these studies has been highly inspirational, in the design of automated steering controllers.

3 Modeling

The dynamics of a passenger vehicle can be described by a detailed 6-DOF nonlinear model [Pen92]. 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 3) is described by the following variables and parameters:

- \mathbf{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:

$$F_f = c_f \alpha_f$$

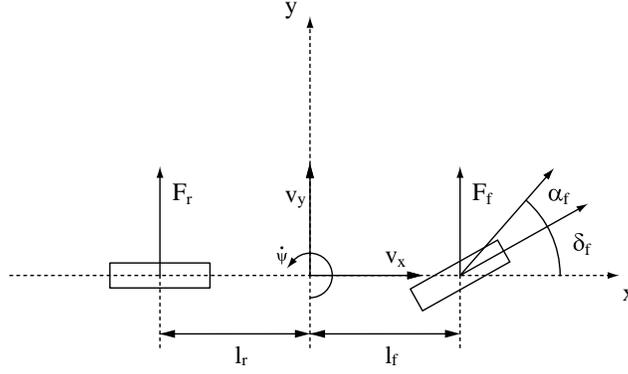


Figure 1: The motion of the vehicle is characterized by its velocity $\mathbf{v} = (v_x, v_y)$ expressed in the vehicle's inertial frame of reference and its yaw rate $\dot{\psi}$. The forces acting on the front and rear wheels are F_f and F_r , respectively.

$$F_r = c_r \alpha_r \quad (1)$$

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:

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

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

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

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

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

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

3.1 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:

$$\dot{y}_L = v_x \varepsilon_L - v_y - \dot{\psi} L \quad (5)$$

$$\dot{\varepsilon}_L = v_x K_L - \dot{\psi} \quad (6)$$

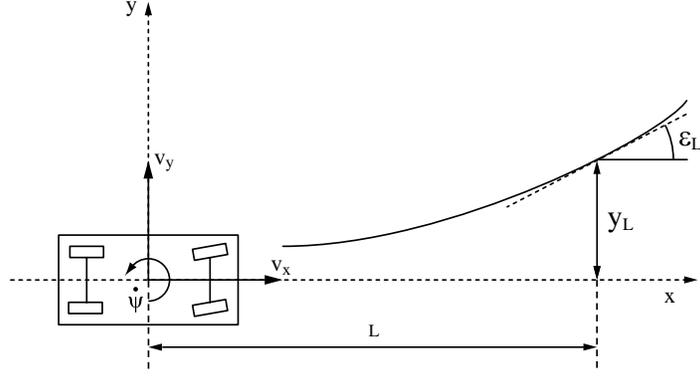


Figure 2: The vision system estimates the offset from the centerline y_L and the angle between the road tangent and heading of the vehicle ε_L at some lookahead distance L .

Where K_L represents the curvature of the road.

3.2 Combined Model

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

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

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

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

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

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

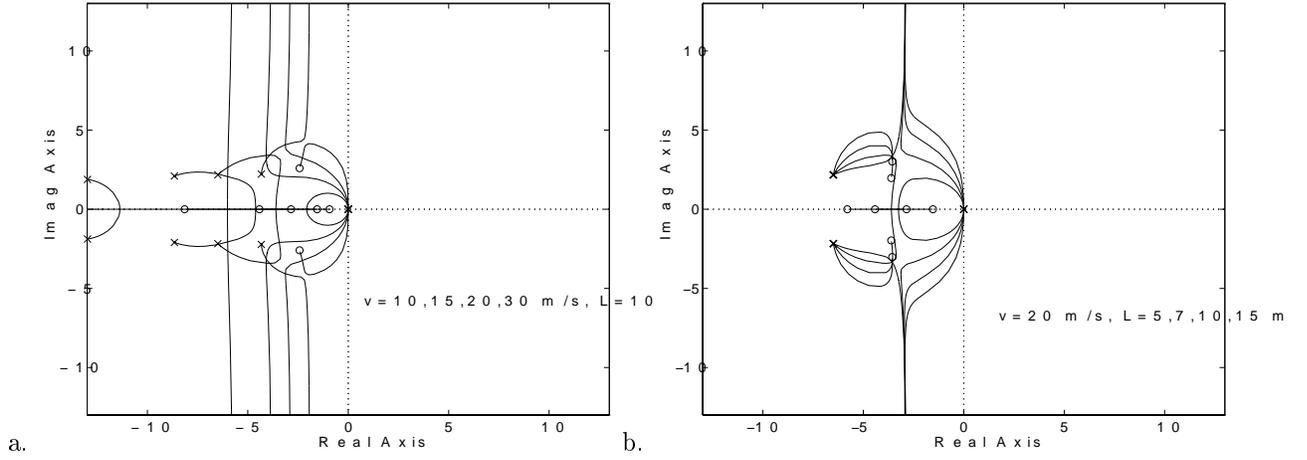


Figure 3: (a) Root locus of $V(s)$ for velocities $v_x = 10, 15, 20, 30$ m/s and fixed look-ahead distance $L = 10$ m. Increasing the velocity v_x moves both the poles and zeros towards the imaginary axis. (b) Increasing the look-ahead distance L moves the zeros of the transfer function closer to the 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 $V(s)$.

3.3 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:

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

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 30 m/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 $V(s)$.

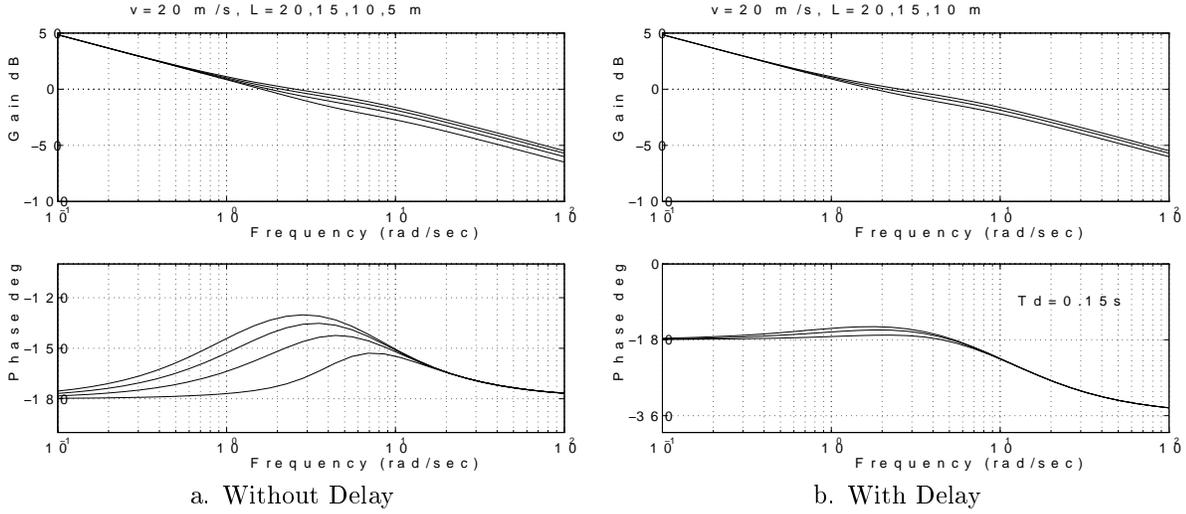


Figure 4: (a) Bode plot $V(s)$ for varying look-ahead $L = 5, 10, 15, 20\text{m}$ at $v = 20\text{m/s}$ with no delay. Increasing the look-ahead adds substantial phase lead at the crossover frequency. (b) The presence of the delay adds an additional phase lag over the whole range of frequencies. The 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.

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 $D(s) = e^{-T_d s}$ which is placed in series with the vehicles transfer function $V(s)$. The processing delay T_d in our implementation was 57 milliseconds.

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 $40\text{-}60^\circ$ phase margin at the crossover frequency. Bode diagrams of $V(s)$ and $V(s) D(s)$ 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 $v = 20\text{m/s}$ and $L = 20\text{m}$, the maximum processing delay one can afford to tolerate before bringing the phase margin to zero is $T_{dmax} = 0.39\text{s}$, 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.

4 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



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

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 [Fau93]. 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) .

$$\begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \propto H \begin{pmatrix} u \\ v \\ 1 \end{pmatrix} \quad (10)$$

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 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

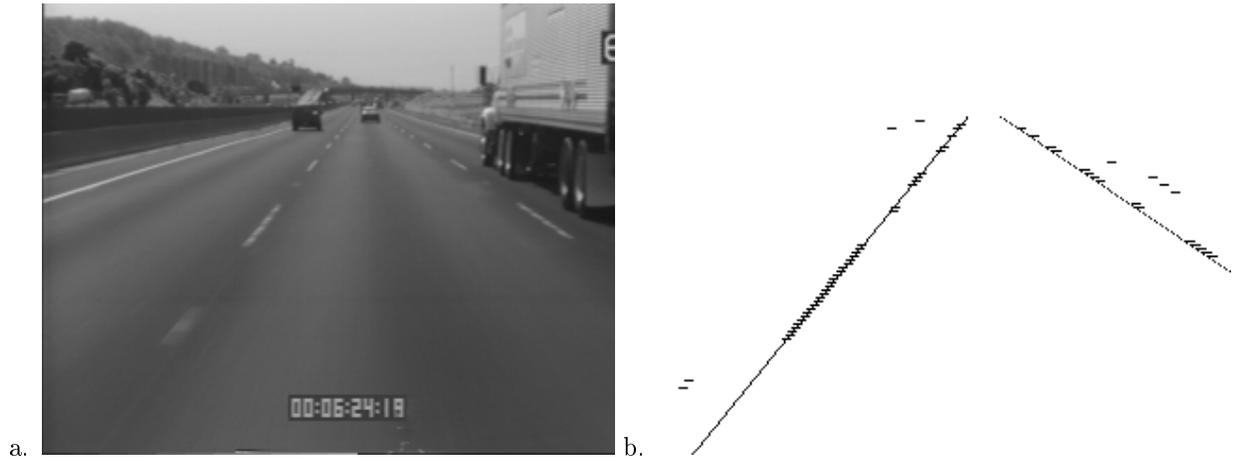


Figure 6: These figures show the performance of the lane extraction system on a typical input image

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 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).

5 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).

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

$$\mathbf{y} = C'\mathbf{x}' \quad (12)$$

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

Note that the state vector \mathbf{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 .

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

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 \mathbf{y}' is used to update an estimate for the state of the system $\hat{\mathbf{x}}'$ as shown in the following equation:

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

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

The gain matrix L can be chosen in a number of ways [Gel74], 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.

6 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 [GTP96]. 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.

6.1 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 $V(s)$ 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:

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

where $C(s)$ 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, $V(s)$, 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.

6.2 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 $\mathbf{u} = K\mathbf{x}$. 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.

6.3 Input-Output Linearization

Input-output linearization is typically used to linearize nonlinear systems by state feedback as described in [Isi89]. 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 vehicle's 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:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) + g(\mathbf{x})\delta_f \quad (16)$$

$$y_L = h(\mathbf{x}) \quad (17)$$

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

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

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

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.

6.4 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 $[\dot{v}_y, \ddot{\psi}, \dot{y}_L, \dot{\epsilon}_L]^T$ to 0.

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

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.

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

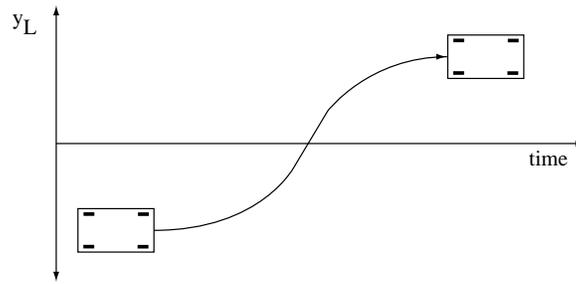


Figure 7: Lane change maneuver

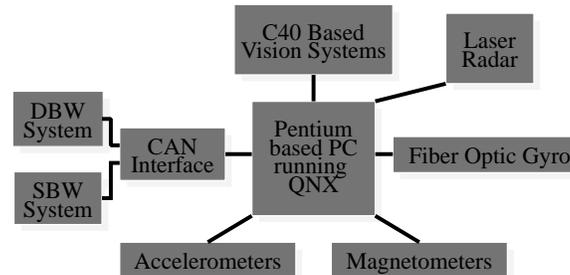


Figure 8: System Diagram

6.5 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.

7 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 of sensors which give it information



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

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.

8 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 10) and our experiments were run at speeds of approximately 75mph to simulate actual 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 10 and 11 describe the performance of tested control strategies without and with the feedforward term.

Figures 10a, 10b and 10c 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 10d, 10e and 10f indicate the velocity profiles during these runs while Figures 10g, 10h and 10i 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/O linearized controller is quite good at lower velocities but the ride becomes a little rougher at higher velocities.

The plots in Figure 11 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 essen-

tially 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).

9 Conclusions

This paper 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.

Acknowledgment. This research has been supported by Honda R&D North America Inc., Honda R&D Company Limited, Japan, PATH MOU257 and MURI program DAAH04-96-1-0341.

References

- [aMH91] D. Raviv and M. Herman. A 'non-reconstruction' approach for road following. In *SPIE proceedings on Intelligent Robots and Computer Vision*, pages 2–12, 1991.
- [BB97] Massimo Bertozzi and Alberto Broggi. Vision-based vehicle guidance. *Computer*, 30(7):49, July 1997.
- [DM92] Ernst D. Dickmanns and Birger D. Mysliwetz. Recursive 3-d road and relative ego-state recognition. *IEEE Trans. Pattern Anal. Machine Intell.*, 14(2):199–213, February 1992.
- [ECR92] B. Epiou, F. Chaumette, and P. Rives. A new approach to visual servoing. *IEEE Trans. on Robotics and Automation*, 8(3):313–326, June 1992.
- [Fau93] Olivier Faugeras. *Three-Dimensional Computer Vision*. MIT Press, 1993.

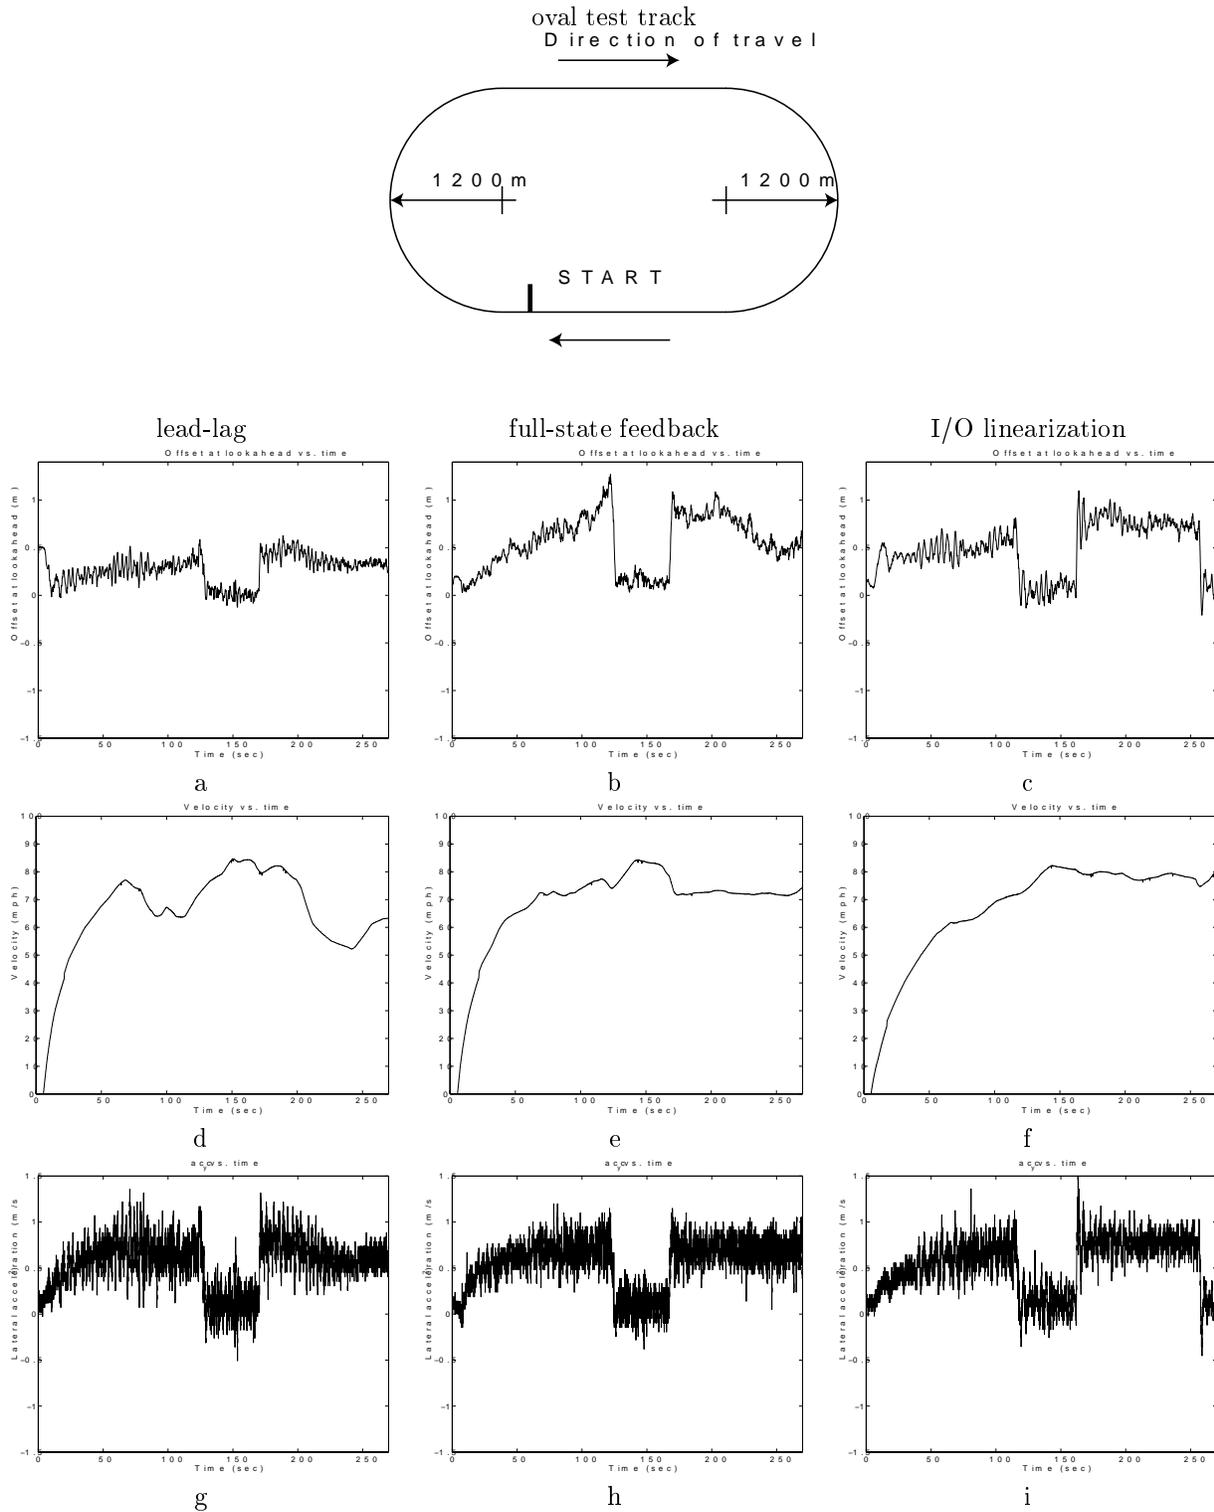


Figure 10: This figure presents a side by side comparison of the results obtained when our test vehicle was driven on an oval track under each of the control schemes that were implemented. The first column of plots correspond to data collected under the lead-lag control scheme, the second to the full-state 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 linearization and 6 meters in the case of full state feedback. Figures d, e and f indicate the velocity profiles used on these runs while Figures g, h and i denote the lateral acceleration experienced at the center of gravity of the vehicle. The lead-lag controllers tracking performance is superior to that of the other two control strategies.

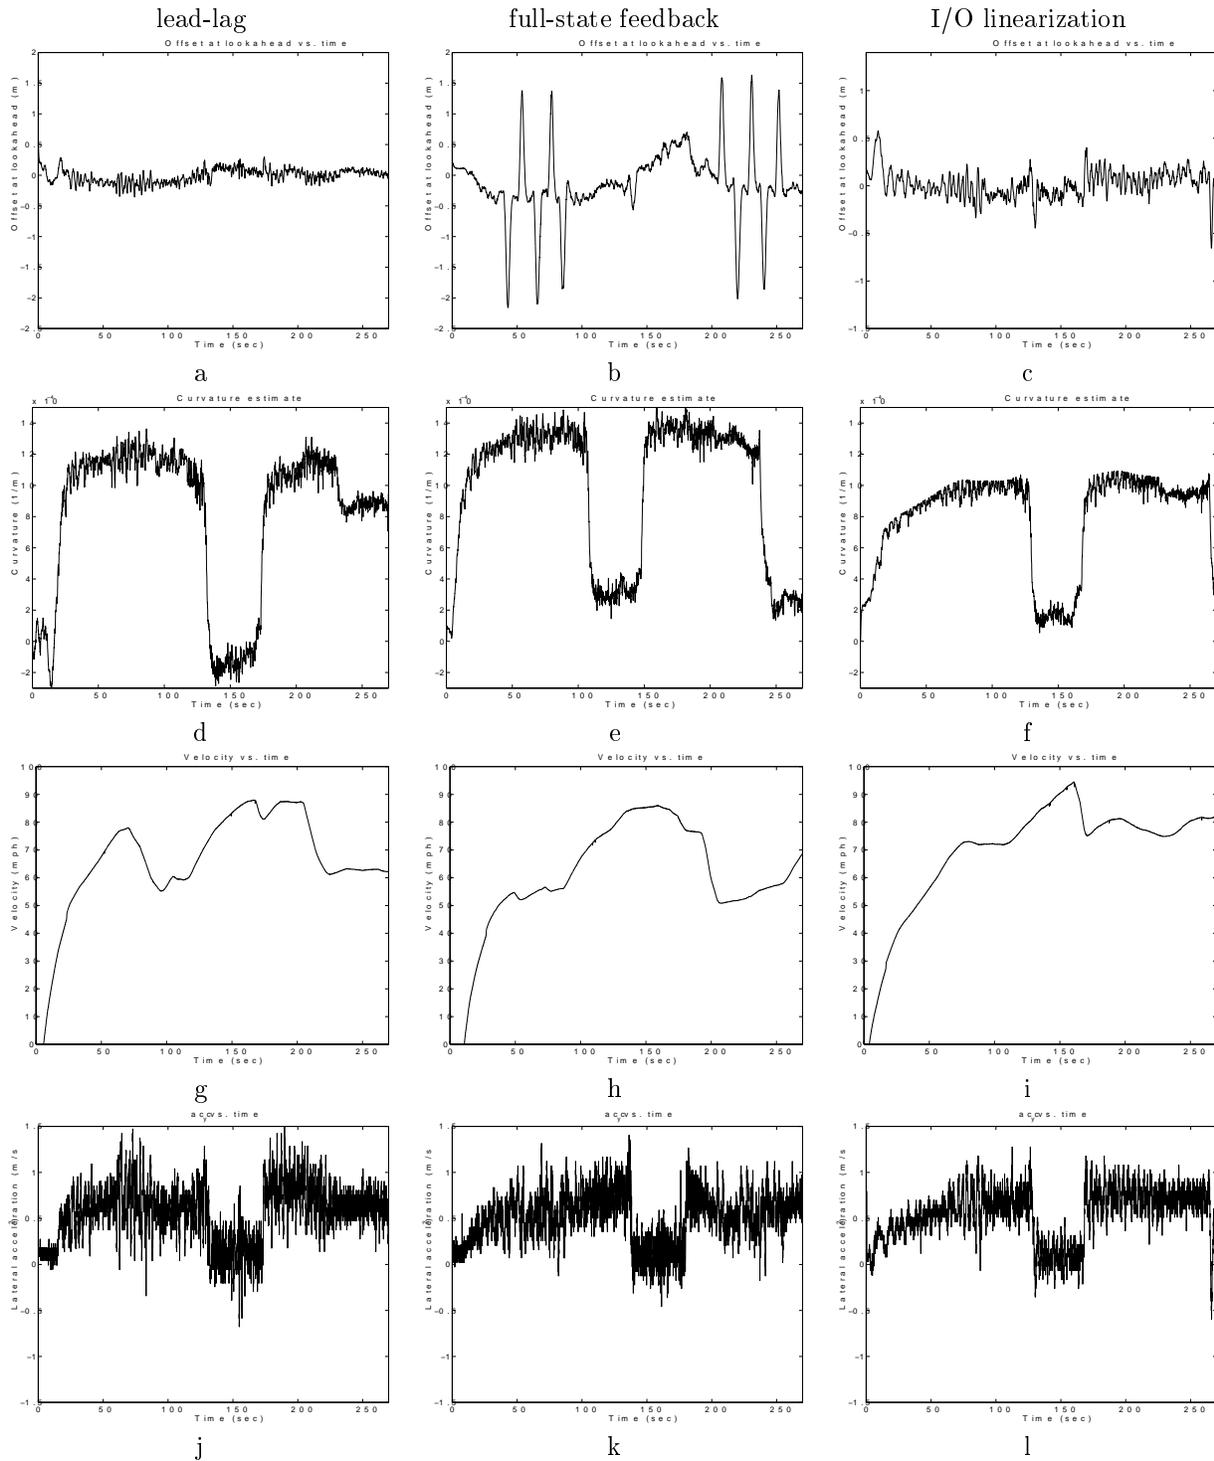


Figure 11: These plots demonstrate the effect of the feedforward control term on the overall tracking performance for all tested control strategies. The first row of plots indicates the tracking performance measured in terms of the offset at the lookahead, the second row depicts the curvature estimate 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 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 the 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).

- [Gel74] Arthur Gelb. *Applied Optimal Estimation*. MIT Press, 1974.
- [GTP96] J. Guldner, H.-S. Tan, and S. Patwardhan. Analysis of automated steering control for highway vehicles with look-down lateral reference systems. *Vehicle System Dynamics*, 1996.
- [Hes90] R. A. Hess. A control theoretic model of driver steering behavior. *IEEE Control Systems Magazine*, pages 3–8, August 1990.
- [HNH⁺97] Martin Herman, Marilyn Nashman, Tsai-Hong Hong, Henry Schneiderman, David Coombs, Gin-Shu Young, Daniel Raviv, and Albert J. Wavering. Minimalist vision for navigation. In Y. Aloimonos, editor, *Visual Navigation, From Biological Systems to Unmanned Ground Vehicles*. Lawrence Erlbaum Associates, Mahwah, New Jersey, 1997.
- [Isi89] Alberto Isidori. *Nonlinear Control Systems*. Springer Verlag, 1989.
- [Lan96] M. F. Land. The time it takes to process visual information when steering a vehicle. In *ARVO Poster B248, Investigative Ophthalmology*, 1996.
- [LH95] M. Land and J. Horwood. Which parts of the road guide steering. *Nature*, 377(28), September 1995.
- [LL94] M. F. Land and D. N. Lee. Where do we look when we steer? *Nature*, 369(30), June 1994.
- [ÖÜH95] 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.
- [ÖÜH97] Ü. Özgüner, K. A. Ünyelioglu, and C. Hatipoğlu. Steering and lane change: A working system. In *IEEE Conference on Intelligent Transportation Systems*, 1997.
- [Pen92] H. Peng. *Vehicle Lateral Control for Highway Automation*. PhD thesis, Department of Mechanical Engineering, U.C. Berkeley, 1992.
- [Pom95] Dean Pomerleau. Ralph: Rapidly adapting lateral position handler. In *Proceedings Intelligent Vehicles 1995*, pages 54–59, 1995.
- [THKS88] Chuck E. Thorpe, Martial Hebert, Takeo Kanade, and Steve Shafer. Vision and navigation for the Carnegie-Mellon navlab. *IEEE Trans. Pattern Anal. Machine Intell.*, 10(3):362–373, May 1988.