

Predicting Lane Position for Roadway Departure Prevention

Parag H. Batavia Dean A. Pomerleau Charles E. Thorpe
{parag|pomerlea|cet@ri.cmu.edu}
Carnegie Mellon University - Robotics Institute
Smith Hall 200
Pittsburgh, PA, 15213, USA

Abstract-- In this paper, we describe a preliminary analysis of driver data collected during a recently completed small scale data collection effort. We will demonstrate that a popular method for computing Time To Lane Crossing (TLC) does not always accurately predict the driver's actual TLC. We will then use a memory based learning approach to show why this is. Finally, we will present results in predicting the driver's future lane position using the new memory based approach.

Index terms-- lane departure warning, driver model, Time to Lane Crossing, TLC, memory based learning, Navlab

I. INTRODUCTION

This paper describes recent work in analyzing driver data collected using a CMU Navlab, along with work in developing driver models for use in single vehicle roadway departure (SVRD) prevention systems. This problem has significance due to the large number traffic fatalities which occur annually due to SVRD.

Current warning systems, such as Pomerleau's RALPH warning system [8] use physics based approaches, such as looking at the position of the vehicle in the lane, or looking at the direction the vehicle is pointing relative to the lane and calculating a time to lane crossing (TLC) [5] metric, which is the time it would take for the first tire of vehicle to cross a lane boundary. While these types of systems work well, we believe that more intelligent modeling of the driver's behavior could result in an increase in warning time and a decrease in false alarm rate.

Towards this end, this paper looks at two different modeling approaches. A kinematic approach is used as a baseline, while a memory based learning approach is investigated for improved performance. These models are used to predict the future lane position of a driver, given current lane position and lateral velocity.

II. MOTIVATION

In 1996, there were over 37,000 fatal automobile accidents, in which 42,000 people were killed. While

there are many different causes of accidents, those that involve a single vehicle are frequently caused by driver inattention or incapacitation, leading to roadway departure. Of the 37,000 fatal accidents in 1996, over 21,000 were single vehicle accidents. These 21,000 accidents resulted in 22,500 fatalities, or 56% of the total [10]. The combined cost of *all* accidents is estimated to be over \$150 billion per year.

Lane departure warning systems can be used to either help prevent these accidents, or reduce their severity by providing advanced warning for the driver to initiate corrective action. We believe that developing effective driver models will lead to increased reliability and acceptability of SVRD warning systems.

III. PREVIOUS WORK

While the literature for general driver modeling goes back over 40 years, work in modeling for SVRD prevention is more recent.

The RALPH lane tracker and warning system [8] uses TLC to determine when the driver is in danger of crossing a lane boundary. The TLC estimate is calculated using instantaneous lane position and lateral velocity. An allowance is made for curve-cutting behavior, in which drivers tend to shift towards the inside of a curve. While the system works well, the TLC threshold has to be kept quite low (0 seconds, or when a wheel touches a lane boundary) to keep false alarms low. TLC also has a problem in situations where the driver tends to drive very close to the lane boundary. In these cases, the false alarm rate can be high, as small perturbations in driver position can have large effects on TLC.

The CAPC system [4] from the University of Michigan Transportation Research Institute also uses a TLC approach. However, their calculation of TLC is more sophisticated than RALPH's. They use upcoming road curvature, along with vehicle dynamics to predict TLC. While their TLC is presumably quite accurate, they most likely would have the same problems that

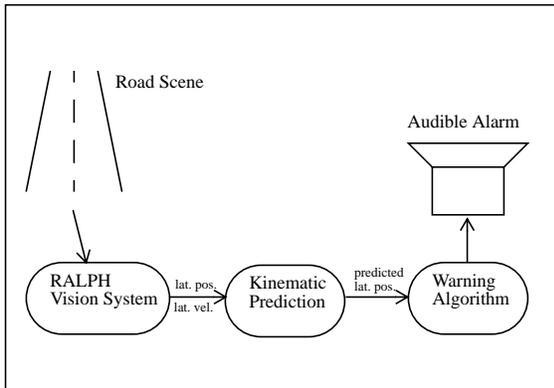


Figure 1. RALPH warning system block diagram.

RALPH does near lane boundaries. However, quantitative data is not available to demonstrate this.

Liu and Pentland [7] have done interesting work using multiple Kalman Filter (KF) models to predict driver action during distinct driver behaviors such as lane keeping, overtaking, and lane changing, and a Hidden Markov Model (HMM) to select among the KF models. Their results indicate that in simulation, different behaviors can be detected with approximately 90% accuracy 1.5 seconds after the beginning of the maneuver.

Others have also done work on driver modeling and vehicle state prediction for SVRD and drowsy driver detection, such as Knippling and Wierwille [6], who did a comprehensive experiment using regression on different vehicle state inputs to detect drowsiness. Isomoto [Isomoto95] et al. look at current and future lateral displacement to determine whether the driver is about to depart the lane. Brattoli et al. [3] also look at lateral displacement to detect lane departures. Batavia [2] used a neural net to predict driver steering wheel action, to model individual drivers and detect uncorrected deviations.

IV. MODEL DESCRIPTIONS

In our experiments, two models were tested, using data and experimental methodology described in Section V.

Our approach to designing an SVRD warning system separates the driver model and vehicle state prediction from the warning algorithm. This has two advantages: It is easier to test separate prediction and warning algorithms when they are not tied together, and the effects of differing levels of complexity for each algorithm can be studied. The RALPH system implements the simplest form of this approach. It uses kinematic prediction to predict future vehicle state, and a warning based purely on lane position. Figure 1 shows a block diagram of RALPH. This paper, however, will only concentrate on experiments on improving the driver model portion.

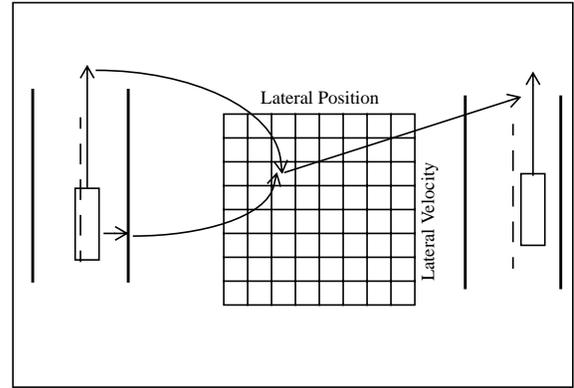


Figure 2. Memory Based Learning approach. The current lane position and lat. vel. is used to look up the most likely future lane position based on previous training data.

A. Kinematic Prediction

Kinematic prediction of vehicle position is a very simple model, which works exactly the way it sounds. Given the current vehicle state, which consists of lane position and lateral velocity, compute what the lateral position of the vehicle will be in t seconds, assuming constant lateral velocity, using $lp' = lp + (t \times lv)$, where lp' is predicted lane position, lp is current lane position, lv is lateral velocity, and t is the prediction time step.

This prediction can be used as input to the TLC warning algorithm. In TLC, the current lane position and lateral velocity is used to determine the time remaining before one wheel touches a lane boundary. If this time falls below n seconds (where n is referred to as the TLC threshold), an alarm is sounded. Equivalently, we can use kinematic prediction to project where we will be in n seconds, and see if our predicted position violates a lane boundary.

B. Memory-Based Learning

Instead of using kinematic prediction as input to TLC, we tested a more complicated model, one which takes driver actions into account to generate a more accurate prediction of future lane position. This model is a simplified memory-based learning (MBL) approach [1].

The training data input state space is current lateral position and velocity, and the output is the actual lateral position t seconds in the future. In this model, all the training data is stored in a 2-dimensional array, where the indices of the array represent a discretized state space, as depicted in Figure 2. Currently, we use a resolution of 0.05m and 0.05m/s for lateral position and lateral velocity, respectively. As the ranges of lateral velocity and position are on the order of +/- 1m and +/- 1m/s, the array size is not large.

During training, each training point is binned into the appropriate area in the memory array. The actual future lane position is added to a list contained in the array loca-

tion. After all the training points are processed, the mean, variance, and mode of each location is computed. To compute the mode, we use an estimation method for continuous data described in [9].

During the recall, or testing phase, the appropriate array element for each query point is found, based on its lateral position and lateral velocity. The output is the mode of all the points which were binned into that element during training. The use of mode instead of mean is explained in Section VI. If there are fewer than 5 sample in the bin, we use the mean, rather than the mode. If there are no samples in the bin (i.e., our training set did not contain any instances of this particular lateral position/velocity pair), we use kinematic projection.

V. EXPERIMENTAL METHODOLOGY

The following sections will look at details of the available data, and the two methods used to predict vehicle state. The explanations are geared towards addressing the following two questions: 1) Given state information (where state is lateral position and lateral velocity), how well can we predict vehicle lateral position at some time t in the future? 2) What combinations of lateral position and lateral velocity are hard to model?

A. Data

We ran our experiments on data collected from 5 drivers who drove Navlab 8, an Oldsmobile Silhouette, on a round trip from Pittsburgh, PA, to Grove City, PA, which is a total distance of approximately 90 miles of highway driving. The full data collection effort included 20 drivers, for a total of approximately 40 hours of data and corresponding video of the road. The data was collected using the RALPH lane tracking system, and includes information such as lane position, lateral velocity, road curvature, longitudinal velocity, turn signal status, and obstacle information collated from radar sensors around the vehicle. An experimenter was present with the subject during the runs, and the subjects had not previously driven the vehicle. The subjects were instructed to drive normally, and make lane changes and other maneuvers as they saw fit.

We only use data on the return trip, as we allow the driver to become acclimated to the vehicle on the outbound trip. While there may be a chilling effect on driver behavior due to the presence of an experimenter and the unfamiliar vehicle, a post-run survey indicated that the majority of the drivers didn't feel (or were unwilling to admit they felt) uncomfortable during the run. The survey also indicated that they didn't believe they drove differently due to the above confounding factors. The 5 drivers were selected because they had different lane position means and variances, allowing for analysis of different driving styles.

Table 1: Data Set Statistics

Driver	LatPos Mean (m)	LatPos STDev (m)	# Left LC	# Right LC
driver_2	0.0353	0.3760	8	8
driver_4	-0.3263	0.3744	5	5
driver_10	-0.1321	0.3831	7	7
driver_14	0.0552	0.2778	4	4
driver_18	0.0364	0.4116	9	9

B. Methodology

The experimental data (of the return trip only) collected from each of the five selected subjects was partitioned into a 10 minute training set and a 10 minute test set. The partitioning was done such that each set had a roughly equivalent population of curves and lane changes. Statistics for these drivers are shown in Table 1, and include lane position mean and standard deviation (over the entire run), and number of left and right lane changes across the training and test set. As nearly all the data was collected on two-lane highway, it isn't surprising that the number of left lane changes equals the number of right.

Due to the nature of the models, only the test set was used for the kinematic method. However, both the training and test sets were used for the memory based learning method, as the memory table is built using the training set, and tested using the test set.

For both models, we used the data to predict what the lateral position of the vehicle would be t seconds in the future, where $t = 0.0$ to 3.0 seconds, in 0.25 second increments. The prediction was compared to the actual lane position data t seconds in the future, and a mean absolute error metric was computed.

VI. RESULTS AND ANALYSIS

A. Kinematic Prediction/TLC Problems

The graph in Figure 3 shows the results of kinematic prediction on 5 drivers, where the prediction time step are as stated above. The mean of the 1-second prediction errors 0.12m . This is not a bad result. However, the raw error rates do not matter as much as where the errors occur.

Figure 4 demonstrates the problem with kinematic prediction. This graph shows 50 seconds of predicted lane position for driver_10. The blue (or darker) graph is a prediction which is generated by looking ahead in the data. The red (lighter) graph is predicted lane position generated using kinematics. In all cases, negative distances are left of lane center, and positive distances are right of lane center. Overall, the two match well. However, the mis-matches occur at inflection points, i.e.,

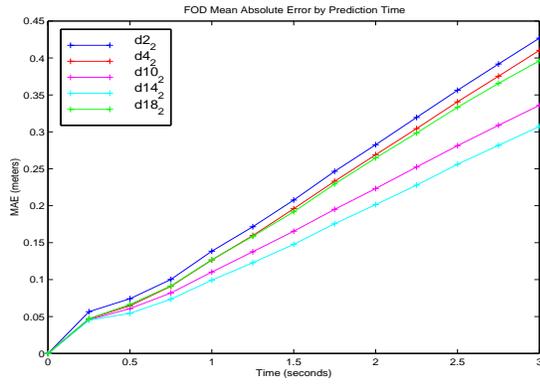


Figure 3. Mean absolute prediction error for 5 drivers using kinematic prediction. The prediction times are from 0.25s to 3.0s.

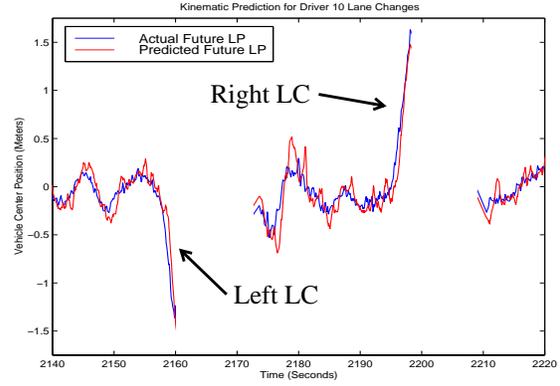


Figure 5. 2 Lane changes of driver_10, predicted using kinematics.

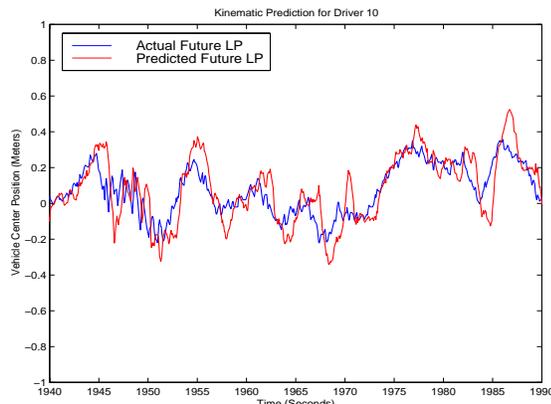


Figure 4. Prediction results for 50 seconds of driver_10 using kinematic prediction. The blue is actual future lane position, the red is predicted.

where the driver reverses direction. This is understandable, as only instantaneous lane position and lateral velocity are used. If a driver is to the left of center, and heading towards the left, a 1st order kinematic model cannot predict that he is likely to correct his trajectory to avoid running off the road.

Predicting vehicle state during normal driving is useful. However, an ideal model would also properly predict future vehicle state in less common cases, such as lane changes. Figure 5 shows the results of using kinematic prediction to predict lane position during two consecutive lane changes. As in Figure 4, the blue (darker) graph shows future lane position, and the red (lighter) graph is predicted future lane position. There is a gap in the data, as when changing lanes, the RALPH vision system can take up to 5-7 seconds to lock onto the new lane. During this time, the lane position estimates are untrustworthy. The overall match between actual and predicted future lane position is quite good. This is understandable, as most smooth lane changes exhibit fairly constant lateral velocity, so kinematic prediction would do well.

The spikes in Figure 4, however, illustrate the potential problem with using a kinematic model and position based

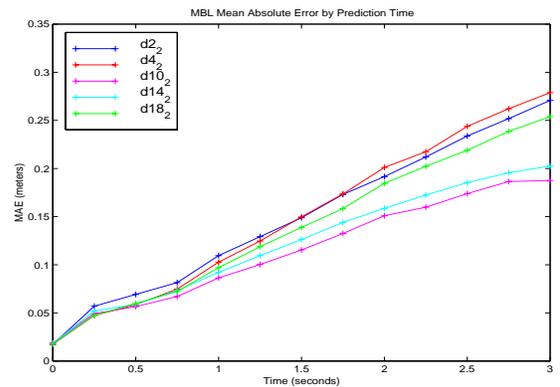


Figure 6. Mean absolute prediction error for 5 drivers. The MBL prediction times are from 0.25s to 3.0s.

warning system with relatively high (≥ 0.5 second) prediction times. An SVRD warning system which uses this method to predict when the vehicle would cross a lane boundary would have high false alarms, as warnings would be triggered in cases where the driver would normally correct his drift. This problem is aggravated by drivers whose mean lane position is significantly away from the lane center. A driver who normally drives to the left of center would be allowed very little leftward drift.

B. MBL Prediction

Using the procedure described in Section V, we computed predictions at varying time steps using the same data as used for the kinematic experiments. Figure 6 shows prediction errors over the same 5 drivers and times as Figure 3. The errors indicate that the MBL prediction is considerably more accurate than kinematic prediction. For instance, the mean of the 1-second prediction errors is 0.097m, which is better than kinematic prediction. However, as we will soon see, these numbers are somewhat misleading.

Figure 7 shows the prediction results for 50 seconds of driver_10. This is over the same data as in Figure 4. The prediction is noisier, as we are currently not performing

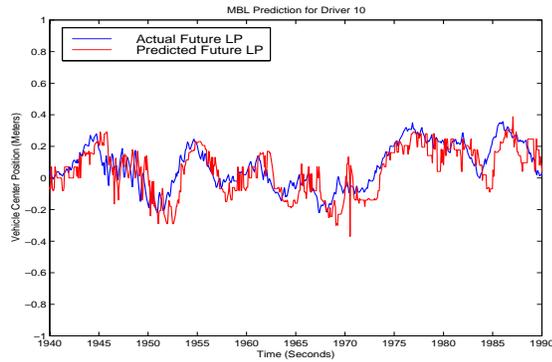


Figure 7. Prediction results for 50 seconds of driver₁₀ using kinematic prediction. The blue is actual future lane position, the red is predicted future lane position.

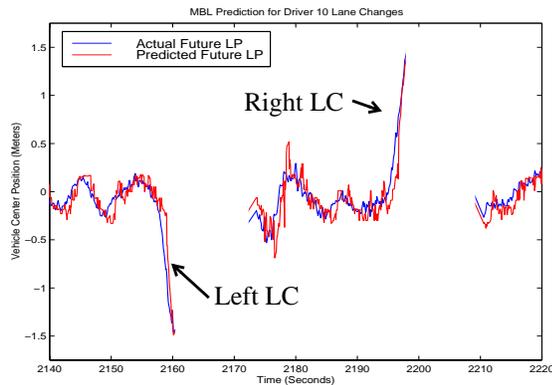


Figure 8. 2 Lane changes for driver₁₀, predicted using the MBL method.

any smoothing over the memory table. Some preliminary experiments in fitting 2nd and 3rd order surfaces to the actual future lane position mode information in the table resulted in smoother predictions, along with improved predictions in areas of the state space where the data was sparse. However, there were instabilities in predictions at the extremes of the state space, as occur during lane changes. What is worth noting is an overall decrease in the magnitude of the peaks that occur at the lane position inflection points. This indicates that the table is better able to predict when the driver is going to change his direction of travel. This, in turn, could lead to a lower false alarm rate. However, spikes can still occur at inflection points if previous training data cause the system to expect a reversal when one does not actually occur in the testing data.

Figure 8 shows the prediction results for the same lane changes as in Figure 5. Again, the actual and predicted future lane position match well, indicating that this approach could model the driver in different situations, such as lane changes.

However, looking closely at Figure 7 reveals something interesting. The predicted future lane position occasionally lags the actual future lane position. Upon closer

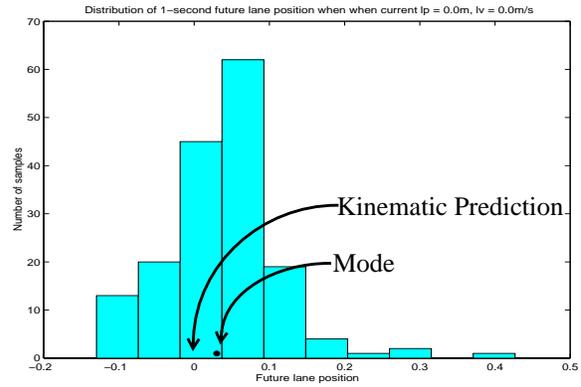


Figure 9. Future lane position distribution for driver₁₀, $lp = 0.0m$, $lv = 0.0m/s$.

inspection, we realized that the MBL method was, in fact, generating predictions that were very close to the current lane position. In other words, if a vehicle is at position p at time t , its most likely position at time $t+1$ is also p ! However, this lag does not appear nearly as strongly during the lane change segments in Figure 8. The reasons for this, and possible solutions, are discussed in the next section.

C. MBL Analysis

To determine why predicted future lane position was similar to current lane position, we need to look at the distribution of predictions inside the memory table. What does the distribution of future lane position look like for a given lane position and lateral velocity?

In the following figures, all distributions are 1-second in the future. This time step was chosen for further analysis because current TLC based warning systems have problems when their threshold is set to 1.0 second. Figure 9 shows a simple case, where the current lane position is 0.0m, and lateral velocity is 0.0m/s. In this case, the mean and the mode of the distribution are both close to 0.03m, and the mode is indicated in the plot. This is not surprising, as when the vehicle is centered, with no lateral velocity, its position one second in the future should be very close to the lane center. This particular driver prefers the left, as shown by the denser sampling in that direction.

Figure 10 shows the 1-second future lane position of a driver who is off to the left by 0.1 meter, and going left at

0.15m/s. The mode of his distribution is -0.15m, which undershoots what kinematics would predict. This is most likely because the vehicle's lateral velocity is not constant during the prediction interval. In this case, kinematic prediction would fail, as it would predict the vehicle's position to be -0.25m. The spread of the distribution is due to the differences in lateral velocity within the samples in the memory table bin centered at -0.1m and -0.15m/s. This indicates that lateral acceleration would be useful to

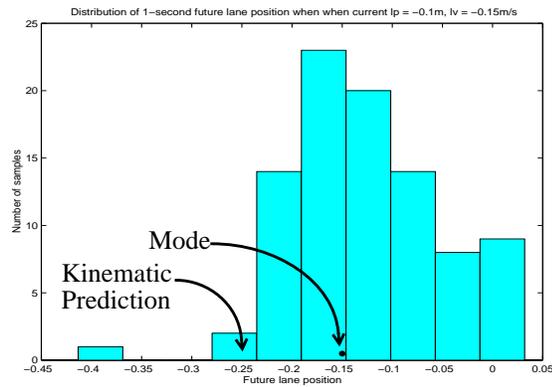


Figure 10. Future lane position distribution for driver 10, $lp = -0.1m$, $lv = -0.15m/s$.

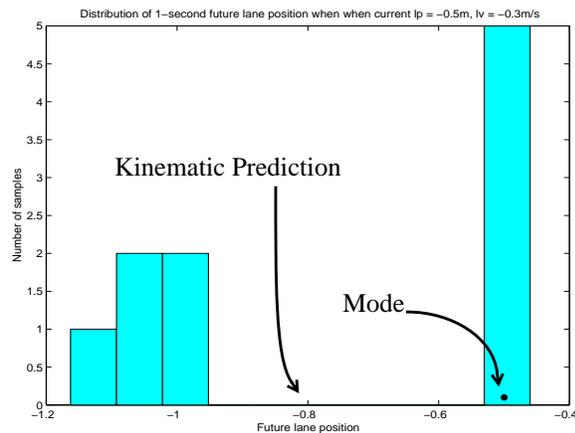


Figure 11. Future lane position distribution for driver 10, $lp = -0.5m$, $lv = -0.3m/s$

include, both to increase the accuracy of kinematic prediction, and to add as a 3rd dimension to the MBL table.

Finally, Figure 11 shows another problem with not including acceleration information. This distribution is the 1-second prediction of driver 10 when he is 0.5m to the left of lane center, and going left at 0.3m/s. The distribution is distinctly bi-modal. In about half the cases, the vehicle is at 0.5m to the left after 1-second, which is where he started from. This indicates a correction maneuver took place to keep the vehicle from drifting off the lane. The other times, the vehicle continued to the left, indicating a lane change. It was this discovery which led us to use the mode of the distribution, rather than the mean.

VII. CONCLUSION AND FUTURE WORK

In this paper, we have presented a preliminary analysis of the performance of two models - kinematic projection, and memory based learning, using a significant amount of real world data. Both models made use of only instantaneous lane position and lateral velocity. We demonstrated the limits of kinematic projection in regards to predicting where the driver will reverse course. We then showed how a memory based learning approach can ameliorate this problem, but introduces issues of its own, such as how to

properly generate a prediction when the distribution of future lane positions given current lane position and lateral velocity can be non-gaussian.

Future work will involve further statistical analysis of the distributions in the memory table, to determine how to best produce accurate predictions. We will also look at using accelerometer data to improve both kinematic and memory based predictions. It may be that incorporating acceleration data would be enough to separate the bimodal distribution in Figure 11, allowing for a Kalman Filter to be used to generate optimal estimates. However, the incorporation of an accelerometer means that the system would no longer be able to rely solely on vision.

VIII. REFERENCES

- [1] Christopher G. Atkeson, Andrew W. Moore, and Stefan Schaal. Locally weighted learning. *Artificial Intelligence Review*, 1997.
- [2] Parag Batavia. Driver adaptive warning systems. Technical Report CMU-RI-TR-98-07, Carnegie Mellon University, March 1998.
- [3] M. Brattoli, R. Tasca, A. Tomasini, E. Chioffi, D. Gerna, and M. Pasotti. A vision-based off-road alert system. In *Proceedings, Intelligent Vehicles*, 1996.
- [4] Robert D. Ervin et al. The crewman's associate for path control: An automated driving function. Technical Report UMTRI No. 95-35, University of Michigan, 1995.
- [5] Hans Godthelp, Paul Milgram, and Gerard J. Blaauw. The development of a time-related measure to describe driver strategy. *Human Factors*, 26(3):257-268, 1984.
- [6] Ronald R. Knipling and Walter W. Wierwille. Vehicle-based drowsy driver detection: Current status and future prospects. In *Proceedings of IVHS America*, pages 245-256, April 1994.
- [7] Andrew Liu and Alex Pentland. Towards real-time recognition of driver intentions. In *Proceedings, IEEE Intelligent Transportation Systems Conference*, Boston, MA, November 1997. IEEE.
- [8] Dean A. Pomerleau. Ralph: Rapidly adaptive lateral position handler. In *Proceedings, IEEE Symposium on Intelligent Vehicles*, Detroit, MI., Sept. 1995. IEEE.
- [9] William H. Press, Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C*. Cambridge University Press, 1st edition, 1990.
- [10] National Highway Transportation and Safety Administration. Fatality analysis reporting system. On the web at <http://www-fars.nhtsa.dot.gov/fars>.