Driver-Adaptive Lane Departure Warning Systems

Parag H. Batavia

September 20, 1999

CMU-RI-TR-99-25

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania, 15213

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy

This work was supported in part by the National Science Foundation (NSF) under a graduate research fellowship, by the National Highway Traffic Safety Administration (NHTSA) under contract DTNH22-93-C-07023, and by the Federal Highway Administration (FHWA) under cooperative agreement DTFH61-94-X-00001 as part of the National Automated Highway Systems Consortium.

Acknowledgements

First, I would like to acknowledge my parents and sister, who have supported me emotionally and financially for the last four years. Next, my advisors, Charles E. Thorpe, and Dean A. Pomerleau, who have helped me grow as a researcher, and have always provided encouragement, good ideas, and interesting stories. The other members of my committee, Illah Nourbakhsh and George Bekey, have been invaluable in the creation of this work. I would like to especially thank Illah for allowing me to take up many hours of his time, and for introducing me to many interesting projects. George Bekey has been a part of my academic life now for eight years, and has always been supportive. Thank you. I would like to thank/blame Jeff Trendel, of Oak Creek High School, for getting me started on this whole robotics kick, by making me fix Heathkit robots over and over again.

I have to thank the group at the USC Robotics lab for taking me on as an undergrad, and giving me the tools necessary to succeed here: Andrew Fagg, Gaurav Sukhatme, Michael McHenry, Jim Montgomery, and others. In particular, I would like to thank Andy for forcing me to come to CMU. :)

I would like to thank my many officemates over the years, who have provided ideas, conversation, and goof-off time. In no particular order: Kevin Lenzo, Kristie Seymore, Chris Leger, Mei Han, Jeff Smith, Diana Labelle, Henry Schneiderman, and John Hancock. Special thanks to John for putting up with my “frustrated months”. I’d include other friends in this category, but there are too many to list, and they know who they are.

Finally, thanks to the Robotics Institute community for providing an interesting, nourishing place to conduct research and form friendships. I’ll miss this place!

Oh yes, one *last* thank you. Thank you to *all* the test subjects who were kind enough to collect data for me. Confidentiality prevents me from mentioning your names, but you made this thesis possible.

Abstract

Each year, there are thousands of car accidents in the U.S. These accidents claim many lives, and cost billions of dollars. There are many different types of accidents, including rear end collisions, side swipes, head on collisions, collisions with static obstacles, accidents while merging or changing lanes, and driving off the road. Mandatory seat belt usage, air bags, lower speed limits, rumble strips, and stricter vehicle safety requirements have all helped to reduce the number of accidents and fatalities. However, it is now possible to do more, by using intelligent driver assistant systems.

In this thesis, I concentrate on a particular type of accident, known as Run-Off-Road (ROR). An ROR crash occurs when a single vehicle departs the road, due to either driver inattention, drowsiness, or other incapacitation, and then impacts something, such as a tree or a house. Previous work in preventing ROR accidents mostly makes use of Lane Departure Warning Systems, which are usually vision-based lane trackers. These systems predict when the driver is in danger of departing the road, and trigger an alarm to warn the driver. To warn drivers early enough so they have time to react means that often, alarms are generated in situations where there is no real danger of a crash. These alarms are called nuisance alarms.

My goal is to reduce the number of nuisance alarms, while maintaining adequate time for the driver to respond to a truly dangerous situation. Using real world driving data, I show that achieving this goal requires more intelligent modelling of the driver's behavior than most current systems are capable of. This modelling comes in two forms: A novel "alarm decision model," which takes into account road geometry and past driver behavior, and a training algorithm which tunes certain model parameters to an individual driver.

This new model reduces nuisance alarms, while maintaining adequate warning time, for "loose" drivers who weave excessively. The improvement for "tighter" drivers is less, as current warning systems already do a good job on them. I analyze the reason for improved warning system performance using a memory based learning framework, and show why "loose" drivers are helped more than "tight" drivers.

Contents

CHAPTER 1 Introduction

1.1 Lane Departure Warning Systems

Traffic crashes are a leading cause of fatalities in the United States. Crashes occur in many different settings, and for many different reasons. In 1996, there were over 37,000 automobile crashes involving fatalities, in which 42,000 people were killed [69]. The combined cost of *all* automobile crashes is estimated to be over $150 billion a year. A short, incomplete listing of the different ways people manage to get into crashes include rear end collisions, side swipes, head on collisions, collisions with static obstacles, crashes while merging or changing lanes, and driving off the road. Until recently, efforts to prevent crashes or mitigate their effects have been limited to approaches such as mandating seat belt use, forcing auto manufacturers to include air bags, lowering speed limits, and installing rumble strips. While all these methods have had a positive impact on reducing traffic fatalities, it is now possible to do more, through the use of advanced technology.

In this thesis, I present a system which has the potential to reduce fatalities which are caused by a particular type of crash: Run-Off-Road (ROR), which are also known as Single Vehicle Road Departure (SVRD). ROR crashes involve a single vehicle, which departs the road and then impacts something such as a tree or a bridge abutment. The causes of RORs include inattention, intoxication, incapacitation, drowsiness, and unintended steering wheel motions [29], and result in a significant portion of U.S. traffic fatalities. It has recently been shown that people who use cell phones while driving have crashes which are similar to those of drunk drivers [63]. For instance, a driver talking on a cell phone might not notice that the vehicle is slowly drifting off the road, or may not notice an upcoming curve. A failure to properly control the vehicle in such a case could lead to an crash.

Currently, the most common approach to preventing RORs is the use of rumble strips on road shoulders. Rumble strips are areas of grooved pavement, usually placed about 6-15cm [75] beyond the lane boundary. When a vehicle drifts off the road, its tire hits a rumble strip, which vibrates the vehicle, and makes a loud noise, alerting the driver to take corrective action. This type of system “sounds a warning” when the driver is actually *in* a situation which has been pre-defined as dangerous, i.e., when the vehicle’s tire is a set distance past the lane boundary.

An alternative approach, which does not require infrastructure modification, is to use a system which can detect when the driver is in danger of departing the road, and sound an alarm in time for the driver to take corrective action. These types of systems are collectively known as lane departure warning systems (LDWS). A LDWS uses sensors to look at the road. These sensors track features on the road, such as lane markings, and use this information to determine the vehicle’s position on the road. This vehicle state can then be used, in turn, to either warn when the vehicle is in a particular state (similar to rumble strips), or to actually *predict* when the driver is in danger of departing the road, which rumble strips cannot do. Examples of LDWS are presented in Section 3.2.

All the current approaches to lane departure warning systems make one key assumption: that exploiting individual differences between drivers is not helpful. These systems either use physics based models which ignore driver behavior, such as [10],[21],[32],[38],[62],[68], or model “generic” drivers, such as [58]. Most systems assume that the driver behaves the same in all situations, and pay no attention to differences in driver behavior due to changes in road geometry, traffic environment, and changes over time. Similarly, these systems assume that drivers do not change their steering input over a given prediction time. However, real drivers weave and oscillate over the road. This behavior can lead to alarms in situations which are not dangerous, and only serve to annoy the driver.

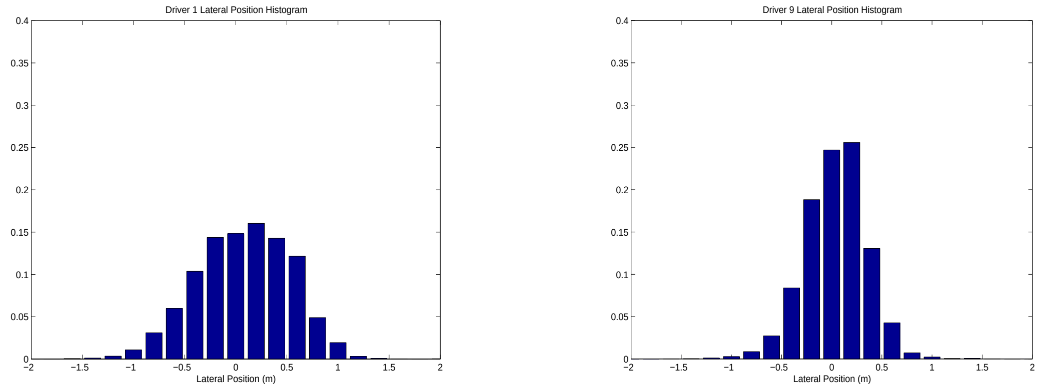

Figure 1-1: Lateral position histograms for a “loose” driver (on the left), and a “tight” driver (on the right).

Intuitively, most people would agree that different people drive differently. We have all seen people weaving down the road. For some, this may indicate an incapacitation, such as fatigue or inebriation. For others, this is simply how they drive. Figure 1-1 shows an example of this. The “loose” driver has a larger spread in lateral position than the “tight” driver, indicating that the “loose” driver weaves more than the “tight” driver. Some people hug one side of the road, as if afraid of getting brushed by traffic on the other side. Others, it seems, sometimes straddle two lanes. Some of these behaviors, such as the tendency to hug one side of the road and corner-cutting, are safe expressions of personal preference. Other behaviors, such as straddling two lanes or weaving wildly, can be unsafe. The difficulty in developing an adaptive lane departure warning system lies partly in adapting to safe patterns in driver behavior, while not adapting to unsafe changes. Ignoring changes and differences in driver behavior can lead to an increase in the number of nuisance alarms (defined below), which could make adoption of a LDWS difficult. Improving warning algorithms and alarm decision models (defined below) can decrease the number of nuisance alarms, increasing acceptance. This leads to my thesis statement, which is:

The number of nuisance alarms can be reduced, while maintaining adequate warning time, using an improved warning algorithm, alarm decision model, and individualized training.

This thesis is the first to tune LDWS performance to specific drivers, and adapt to changes in driver behavior over time. Doing this can result in improved LDWS performance, through the reduction of nuisance alarms, which are alarms in situations which are not dangerous.

1.2 Definitions and Evaluation

First, I define a warning algorithm, and its constituent parts. After that I define difference classes of alarms.

Warning Algorithm: An algorithm which produces an alarm when the driver is in danger of a crash. Most warning algorithms consist of two parts: a vehicle state predictor, and an alarm decision model.

Vehicle State Predictor: A portion of a warning algorithm which takes current vehicle state, and predicts future vehicle state, usually lateral position. There are different ways to accomplish this, including kinematic projection, or data-centric approaches.

Alarm Decision Model: A portion of a warning algorithm which takes predicted future vehicle state and other input, and decides whether to trigger an alarm. Other inputs can include road state (road curvature, shoulder width), and past driver behavior.

Ideally, warnings would only be triggered in situations that would unequivocally lead to a crash if corrective action is not immediately taken. In other words, a warning should be triggered at time *T*, where:

- A warning any later than time *T* would definitely be too late to prevent a crash.
- A warning at time *T* results in the driver successfully avoiding a crash.
- Any crash which is prevented by warning earlier than time *T* could also be prevented by a later warning. The last possible warning time is preferable to limit nuisance alarms.

That is, time *T* is the absolute last moment at which an alarm must be triggered to prevent a crash. It would be nice to be able to determine *T* using a simple rule such as *if vehicle_yaw* > *x*, *lateral_position* < *y*, and *road curvature* > *z*, then we are at time *T*. Unfortunately, it is not possible to define *T* precisely for any given situation. There are many reasons for this; uncertainty about the environment (pavement conditions, shoulder width); uncertainty about the vehicle (tire and brake conditions, slip factors); uncertainty about the driver (reaction time, ability, attentiveness, propensity to oversteer). Beyond these reasons, actually reaching time *T* is quite rare, as in most cases, drivers take corrective action well before *T*. Because it is impossible to determine *T*, drivers need to be warned earlier, to ensure they have enough warning time to react to the situation and avoid a crash.

Therefore, a more practical definition of when it is appropriate to trigger an alarm is required. A simple lane departure, defined as at least one tire crossing the lane boundary, is too stringent. Many drivers, especially truck drivers, tend to touch and cross the lane boundary quite often during normal driving. Sounding a warning in these cases would annoy the driver. Instead, it is possible to define a “substantial lane departure”:

Substantial Lane Departure: A situation in which the vehicle is more than 50% outside the lane. This may or may not result in a crash, depending on road conditions, vehicle conditions, and driver ability.

Substantial lane departures, however, are also quite rare. While it is not unusual for drivers to cross the lane edge slightly, a deviation of almost a meter is uncommon, unless there are extenuating circumstances, like construction zones or debris in the road. It is also possible for lesser deviations to lead to crashes. This leads to my next definition:

ROR Situation (ROR): Any state in which any part of the vehicle departs the lane. Whether an ROR situation leads to a crash depends on the extent of the departure, road state, and driver reaction time and ability.

In the driving datasets which I currently possess, there are no crashes, although there are many ROR situations. In fact, encountering a true ROR crash is, thankfully, very rare, and not something I can wait for. However, to evaluate a warning system, there have to be positive examples of substantial lane deviations. One possibility is to use lane changes, as vehicle state during lane changes can be similar to vehicle state during substantial lane departures due to unintended steering input or inattention, based on the IVHS Countermeasures work [29]. In terms of my thesis statement, from Section 1.1, my goal is to provide as early a warning as possible for lane changes, while minimizing the occurrence of alarms when the driver is not changing lanes. This is the approach I take, which means that this thesis is partly an exploration of lane change behavior. This is further discussed in Section 3.4.3.

Realistically, it is likely that there will be alarm triggers in states which are not lane changes. If these states are substantial excursions, then drivers may feel that it is appropriate for the system to trigger an alarm. As it is still unclear when drivers will or will not consider an alarm appropriate, I present a set of definitions for true, false, and nuisance alarms. If the alarm triggers in a situation which drivers deem normal, they will feel it is a nuisance alarm. Driver perception is therefore very important in classifying an alarm. The definitions of a system designer are irrelevant if the end user does not agree with them.

Hadden et al. [29] have done a simulation study of lane departure countermeasure effectiveness. During analysis of the results, they put forth 6 possible outcomes of a countermeasure intervention. I use this framework, slightly modified, to present the following definitions (Note, when I refer to the driver, I mean a normal driver who is not incapacitated and is paying attention to the road):

Safe True Alarm: An alarm triggered in time to prevent a situation where the driver, in hindsight, recognizes that his actions could have resulted in a crash.

Late True Alarm: An alarm triggered before a crash, but too late for the driver to take corrective action.

Safe Nuisance Alarm: An alarm which triggers in a situation where the driver, in hindsight, does not believe that his actions could have resulted in an ROR crash. Besides the alarm itself, there is no other consequence.

Unsafe Nuisance Alarm: An alarm which triggers in a situation where the driver, in hindsight, does not believe that his actions could have resulted in an ROR crash. In this case, the alarm causes a reaction in the driver which leads to a crash.

False Negative: A situation in which the driver, in hindsight, recognizes that his actions could have resulted in an ROR crash, yet no alarm triggered.

True Negative: A situation in which the driver, in hindsight, does not believe his actions would result in an ROR crash, and no alarm triggered. This is by far the most common outcome.

Burgett [11] mentions another category of nuisance alarms, although I have altered his definition to make it more applicable to this domain:

False Alarm: A nuisance alarm caused by either system perceptual error or an error in the vehicle state prediction or alarm decision model.

The above definitions depend on user perception. The driver data I use in this thesis does not have user reaction to events, so it is difficult to apply the above definitions directly. Therefore, I modify the above definitions for an “offline” environment, when working with pre-recorded data, without access to the subject. These modifications are discussed in Section 3.4.3, and include using lane changes as true alarms, and treating alarms in all non-lane change situations as nuisance alarms.

One question which arises, is how much of an increase in warning time and decrease in nuisance alarm rate is achievable with a modified LDWS? These two issues are very closely related, and improving one generally has a negative effect on the other, as Section 3.4.3.5 shows. System efficiency can be described as being positively correlated with warning time, and negatively correlated with the number of nuisance alarms.

Current systems generally provide enough warning time for drivers to react to an alarm. However, they can suffer from high nuisance alarm rates, particularly for drivers who weave a lot, and frequently drive near or slightly over the lane boundary, which is demonstrated in Section 3.5.3. I show that using an alarm decision model which accounts for more complicated driver behaviors than current systems account for, combined with optimizing the model for a particular driver, can do a better job of predicting dangerous lane departures than standard TLC. This, when applied to a LDWS, can result in warning times which are similar to current systems, with a lower nuisance alarm rate.

1.3 Previous Work

In this thesis, I develop an adaptive warning algorithm which results in a lower nuisance alarm rate. In this section, I present an overview of previous work in both driver modelling and lane departure warning system design. I go into detail about two popular warning methods, Roadside Rumble Strips and TLC, in Section 3.2, where I evaluate their performance using real world data.

1.3.1 Driver Modelling

This section presents previous work in general driver modelling. While this is not the focus of my thesis, this area is large enough and related enough that it is worth briefly reviewing.

One of the earliest works, by Pipes [59], modeled the driver as a gain and a time delay, and modeled the vehicle lateral position as an integration of steering wheel angle. Over the next few decades, that work was expanded upon, as the model of the driver became more complicated and attempted to take into account evidence provided by studies of driver behavior.

Weirwille [73], who has been active in this field for many years, presented an early model which took into account past lateral displacement, future roadway curvature, and driver vantage point. This work showed that information on the upcoming road curvature helps to eliminate the effects of perceptual and reaction lag.

Crossman and Szostak [17] proposed a three level model which combined open loop control of vehicle curvature given upcoming road information, with closed loops around lateral position and lateral velocity. McRuer et al. [46] added a "precognitive" open loop control module, which was used to establish the driver on an appropriate trajectory for lane changes and obstacle avoidance maneuvers. Baxter and Harrison [7] take a previous linear control model, and add a non-linear hysteresis element, in an attempt to model the oscillations of drivers driving on straight roads. Rather than raw vehicle state, they use aim-point error, which is the angle between the vehicle heading and the lane centerline at a certain lookahead distance. Others who have contributed to this field include [1],[5],[6],[7],[12],[13],[19],[20],[22],[23],[26],[28],[30],[36],[37],[42],[45],[64],[66].

The main assumption in control theoretic approaches to driver modelling is that humans, and the vehicles that they control, can be adequately simulated using 2nd order systems. Stochastic and non-linear effects, such as crosswind response, cannot be modeled well using these approaches. Furthermore, it becomes very difficult to take into account environmental effects such as the presence of other vehicles. One area where these approaches have worked well is in car following, as shown by Bekey [8], Chandler [15], Ioannou [31], and Naab [49].

Work in automated vehicle control often makes use of driver models. Neural networks are popular in this domain. CMU's NavLab, which used the ALVINN road follower [61] used a neural net to learn to control a vehicle based on observing a driver, which was combined with a pure pursuit steering model. MacAdam [44], Mecklenburg [47], Cheng [16] and Neusser [52] have also used neural nets to control vehicles. While generally successful in controlling vehicles, only [61] provides any explanation of what the network was learning.

Nechyba [51] and Deng [18] have both investigated the general problem of human skill modelling, and have applied their methods to identifying drivers. Nechyba used a neural net to map vehicle and road state input to steering output, and then validates the models using an HMM based similarity metric that looks at the cross-probability of a sequence of observations generated by both training data, and the model, fed back upon itself. Deng uses an efficient memory based learning approach to compute the log-likelihood of observations, i.e., computing the probability that a given observation (which, in this case, is a PCA-condensed timeseries of vehicle state variables and driver steering command) is likely to be generated by a given driver. While both approaches do well at classifying drivers (based on a subset of the Navlab 8 data described in Section 2.3), they are more properly regarded as system identification methods, rather than explicit driver models. Neither is suitable for on-line prediction or evaluation of driver behavior as part of a lane departure warning system. Zhao [76] has also done work in tracking cars using Kalman filters, which makes use of driver models.

1.3.2 Warning Systems

The Daisy system, developed by Feraric and Onken [54], uses multiple Fuzzy-ART models for different low level situations such as lane changes, passing, cruising, etc., and attempts to learn an individual driver's acceptable levels of Time to Lane Crossing (TLC) and Time to Collision (TTC). A separate model was used to determine what the driver's intent would be in a given situation, to select an appropriate low level model. Heuristic approaches have also been used, where rules are applied to state variables of the type: *If lateral_position > threshold, then sound alarm*, such as the work by Isomoto [32]. Pilutti [58] also used a rule based approach, where the thresholds used by his rules were learned using fuzzy logic.

Takahashi [70] has done some interesting work in developing individualized models of manual gear shifting for engine braking, while driving downhill. Schumann [65] and Tribe [72] have developed warning systems which use active intervention, nudging the steering wheel in the proper direction to avoid a lane departure. There has also been work in predicting the effectiveness of warning systems, by Goodrich [25], Burgett [11] and Pape et al. [55] [56]. A great deal of attention is paid to individual performance and ability in the drowsy driver literature, by Knipling [35], Seigmund [67] and others. There have also been studies to determine where we look when we drive, by Land [39] [40].

With few exceptions, the previous work lacks the following: Performance evaluation on significant amounts of real data, explanations of what is being learned, and what, if any, differences there are between individual drivers which could be exploited by a lane departure warning algorithm. Accomplishing the goal of my thesis requires me to address these issues.

1.4 Thesis Overview

My goal in this thesis is to reduce nuisance alarms by developing a better model for predicting dangerous lane departures, using lane changes as an example of such an event. Another goal is to determine if adapting such a model to an individual driver improves warning system performance.

In Chapter 2, I describe two data collection studies which provided all the data I use in this thesis. The first was begun in 1997, and involved 20 subjects driving Navlab 8 (a minivan with obstacle and lane detection sensors) on a 100 mile round trip. The data collection took about seven months to complete, and was controlled. In all cases, the same vehicle was used, the same route was taken, and an experimenter was present with the subject. Video information of the driver's view was also collected, along with state information such as lateral position, and the status of the various obstacle sensors on Navlab. Although the controlled nature of this study was an advantage in some ways, I was still uncertain whether a more natural setting would provide data which showed greater individual variance. Also, the amount of data per driver was relatively small -- about 1 hour -- which made it difficult to test the learning algorithms I wanted to. Therefore, I did a second data study.

The second study was designed to eliminate some of the confounding factors in the Navlab 8 study. This "naturalistic" data collection effort used a portable AutoTrak lane tracker which was installed in subject's personal vehicles. I hoped that this would capture more natural driving behavior, as the subjects would be driving familiar vehicles without an experimenter present. I selected subjects who were going on long trips (usually greater than 500 miles round trip). Nine subjects volunteered. These nine subjects provided 71 hours of data. Problems in installation and with the data collection system meant that not all of the data which was collected was usable. In the end about 20 hours of clean high quality data were collected, from 5 different drivers.

I use this data to demonstrate three results, which are presented in Chapter 3. The first is an evaluation of current warning system algorithms on realistic data. Most of the work in this area has used driving simulator data, due to the lack of real world data. I evaluate three methods: Roadside Rumble Strips, Time To Lane Crossing (TLC), and AutoTrak, which uses a Future Offset Distance algorithm with hand-tuned parameters. Roadside rumble strips has a low nuisance alarm rate, but a low warning time. Conversely, TLC, as it is normally implemented, has a higher nuisance alarm rate than roadside rumble strips, with a high warning time. The difference in nuisance alarm rate and warning time between roadside rumble strips and TLC depends on the TLC lookahead timestep. AutoTrak has a lower nuisance alarm rate than TLC, with a higher warning time than roadside rumble strips.

The second and third results are presented concurrently. The second result is the development of a more sophisticated warning algorithm and alarm decision model, which is a model that takes future vehicle state and current road state as input, and outputs whether an alarm needs to be triggered. This new alarm decision model takes into account driver behavior such as curve cutting and changes over time, and when used in a Future Offset Distance (FOD) warning algorithm, results in a lower nuisance alarm rate, with a minimal impact on warning onset time. Finally, I present the third result on training FOD parameters for individual drivers. I show that while the benefits of training are debatable for many drivers, drivers who weave a lot, either due to ability, road conditions, vehicle type, or weather (who might be considered “loose” drivers), can be helped by training a driver-specific model.

In Chapter 4, I describe a new way to model and analyze driver behavior. I use a technique called memory based learning (MBL), which is similar to k-nearest neighbor. While most of the previous work analyzes drivers in term of model fits and gross statistics, MBL allows me to look at subtle differences in driver behavior over different parts of the state space. In this case, the state space is lateral position and lateral velocity. Training the MBL table involves iterating over data, and binning it based on its state. When a particular state is binned, the *actual* future lane position (after some lookahead time) is binned. Therefore, we can start asking questions like “When the driver is 10cm to the left, and drifting left at 10cm/s, where is he/she likely to be two seconds later?”. There are a number of interesting observations which come out of this type of analysis. One is that drivers do not always behave according to simple kinematic predictions. Most current warning algorithms (including TLC and FOD) assume that drivers will not change their steering wheel position over the prediction time step. An MBL analysis shows that as one looks further into the future, the actual future lane position in a given state can become bimodal, with one mode indicating a corrective action, and another mode indicating a lane departure or lane change. This bimodality presents a fundamental limitation to warning system performance when using only lateral position and lateral velocity.

I also use probability and information theory to answer why “loose” drivers are helped by individualized training, and why average drivers derive little or no benefit. I show that the training methodology I use changes the shape of alarm/no alarm decision boundary in such a way that parts of the state space which do not necessarily lead to lane changes, but which “loose” drivers occupy more frequently than “tight” drivers, are re-classified as “no trigger” states. This reduces the overall number of alarm triggers, while minimally affecting the warning time for true excursions. Finally, I use entropy measures to show that for certain drivers, the information content in their driving data is lower, and therefore suffer from higher prediction uncertainty than other drivers, which limits the effectiveness of warning algorithms.

Lastly, Chapter 5 presents a summary of the contributions of this thesis, which are methodologies for real world data collection, along with the data collection itself, the Future Warning Distance (FOD) warning algorithm, results in individualized driver modelling and training, and a low-level analysis of the results using a Memory Based Learning (MBL) framework. After that, I present directions for future research. In particular, I discuss the problems created by a lack of data, and suggest the conditions for further data collection studies and field evaluations.

CHAPTER 2 Data Collection

2.1 Introduction

One factor differentiating my thesis from previous lane departure warning work is my use of a large amount of real world driving data. To this end, I have collected over 100 hours of data from 29 different drivers. This data was collected in two different collection efforts, which made use of two different lane tracking systems, with similar performance. In the first study, the subject drove Navlab 8, a mini-van outfitted with a lane tracker. My second study effort involved installing a portable lane tracker in the subject's vehicle.

Most previous efforts on driver modeling and lane departure prevention use data which was collected by having subjects drive simulators [9], [40], [50], [58], [71], [74]. Simulator data has certain advantages: 1) it is free of sensor noise, 2) simulator studies are amenable to controlled situations such as road and weather condition and traffic density, allowing for the observation of only the desired effects, 3) subjects can be asked to perform the experiment again, if needed, as most simulator studies are of relatively short (less than 2 hours) duration, and 4) drivers can be placed in dangerous situations.

In contrast, real world studies in this area usually suffer from (unmodeled, correlated, and non-gaussian) sensor noise, and uncontrolled effects due to weather, road, and traffic conditions. They are also generally one-shot efforts. If a driver has volunteered to have data collected over a 10 hour trip, he/she will not be willing to go on the trip again if an error occurred in data collection. Also, drivers can not be placed in dangerous situations, such as those which would lead to roadway departures.

Despite these limitations, real world data has many advantages. The primary advantage is that if the goal of driver modeling and lane departure prevention work is to build systems which can work in the real world and prevent single vehicle roadway departures, then the models and algorithms used must be validated on real world data. Most of the current work based on simulation studies, however, leaves real world implementation as a future problem to be solved. In doing so, most researchers do not convincingly demonstrate that their methods are applicable in the real world.

Other advantages of real world data are mainly answers to problems with simulators. The fidelity of most driving simulators leaves something to be desired. While there are some high fidelity simulators in existence [29], the majority are fixed based with less than a 180 degrees field of view. The “feel” of driving these types of simulators is unlike driving a real car. Therefore, the data collected from them, particularly when related to human control strategy, is suspect.

This chapter discuses the two data collection studies, including the lane trackers which were used, the methodology, and an overview of the resulting data. This data forms the basis for the lane departure prevention work which I present in upcoming chapters.

2.2 Lane Trackers

I use two lane trackers to collect the data used in this thesis. A lane tracker is a system which consists of a computer and a sensor which can determine the state of the vehicle and the road. These systems provide information such as vehicle lateral position and velocity, lane width, road curvature, and lane marker types (solid or dashed).

Lane trackers can use sensors such as forward or downward looking cameras, lasers, or carrier phase GPS combined with accurate road maps. The lane trackers which I use are based on a forward looking camera mounted on the vehicle’s windshield, looking out at the road ahead. In such systems, the image of the road is analyzed, and the lane markers or other

Figure 2-1: Navlab 8

useful features are detected. The lane tracker uses the position and orientation of these features to determine the vehicle and road states. This information can then be used by a lane departure warning algorithm to warn drivers who are in danger of running off the road.

The first lane tracker, the Rapid Adaptive Lateral Position Handler (RALPH) [62] was used for the Navlab 8 data collection described in Section 2.3. The second system used, AutoTrak, is a more recent version of RALPH optimized for lane departure warning, and was used in the naturalistic data collection effort. The following sections describe the data collection efforts, including the lane trackers used, the methodology, and the resulting data.

2.3 Navlab 8 Data

The goal of this study was to collect driving performance data to evaluate differences in ability and style, while controlling the environment as much as possible, thereby eliminating any confounding influence. Subjects were asked to drive Navlab 8, which is shown in Figure 2-1. Navlab 8 is an Oldsmobile Silhouette mini-van outfitted with a RALPH lane tracking system, along with various other sensors such as radar, laser, and GPS.

State Variable	Range	Units	LSB
Time of Output	0-860,400	seconds	1 msec.
System Confidence	0 - 1	n/a	0.01
Lateral Position	+/- 10	meters	0.01
Average Road Curvature.	+/- 0.1	1/m	0.00001
Vehicle Curvature	+/- 0.1	1/m	0.00001
Vehicle Yaw	+/- 4.0	degrees	0.33
Vehicle Yaw Rate	+/- 30	deg/sec	0.01
Lane Width	2.5 - 5	meters	0.04
Steering Position	encoder counts / steering wheel motion		
Turn Signal State	left or right turn signal active		
Estimated Lane	left, center, or right lane		
Time to Lane Crossing (TLC)	-1.0 - 6.0	seconds	0.03
RALPH Event	lane departure event flag		
GPS X Coordinate	+/- 20M	meters	0.01
GPS Y Coordinate	+/- 20M	meters	0.01

TABLE 2-1 RALPH State Output

The primary reason for using Navlab 8 for this study was that at the time, I did not have a portable lane tracker to install in a subject's vehicle. A secondary reason was to ensure that all subjects would be driving the same automobile to eliminate any variability due to differences in vehicle dynamics or control.

2.3.1 RALPH Lane Tracker

The RALPH system was originally developed for use as a vision based lane tracker for automated lateral control of a vehicle. The system runs on Pentium class PCs, and interfaces to a camera to track the lane boundaries and provide vehicle state information such as lateral position, vehicle yaw, and upcoming road curvature. A partial state output is given in Table 2-1.

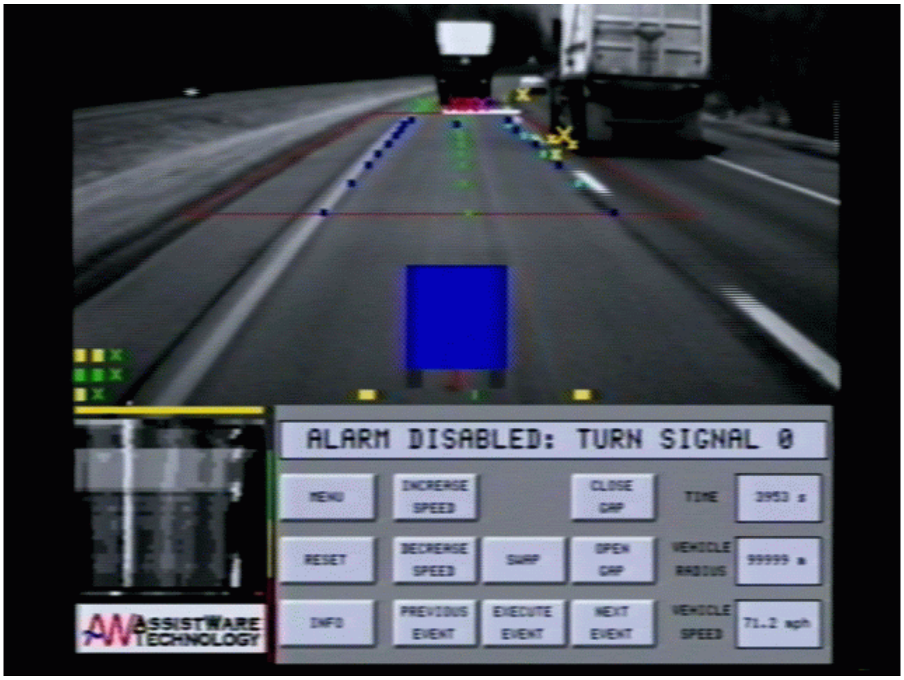

Figure 2-2: RALPH Screen Shot

2.3.1.1 Operation

RALPH operates as follows (Refer to Figure 2-2 for details). The system looks out between 30 and 120m ahead, and a trapezoid covering this area in the image is sampled. The image inside the trapezoid is geometrically warped to produce an overhead view. This transformation causes lane markings on a straight road to become parallel.

A one-dimensional vertical histogram of the image intensity in the transformed area is then computed. This histogram has spikes that designate the location of the bright lane markings. The lane marking positions are computed using the location of the spikes. Upcoming road curvature is then computed using template matching. Templates of the histograms produced by various road curvatures are stored, and compared against the actual histogram. The closest match indicates the curvature of the road.

In Figure 2-2, the trapezoidal area indicates the area being sampled. After the lane markings and curvature are found, they are marked on the image as the dark dots superimposed over the lane markings. The lighter dashed line between the lane markings indicates where RALPH thinks the center of the lane is. The menu in the lower half of the screen is used as part of the automated control system, and allows the driver to change settings such as vehicle speed and following distance, and to initiate actions such as lane changes.

2.3.1.2 Algorithm Characteristics

The lateral position, lateral velocity, and upcoming road curvature estimates generated by RALPH were all accurate enough for active control of NavLab at highway speeds. However, the necessity of using RALPH to control NavLab led to design decisions which reduced the potential accuracy of the lateral position estimates. The primary such decision was to trade off lateral position accuracy for upcoming road curvature accuracy (which is very important when driving at highway speeds) by looking out up to 120m away. This improves curvature accuracy; 100m of lane features over which to compute curvature provides more data than would 20m. However, at 120m away in the image, the lane markings are harder to detect and localize, resulting in reduced lateral position accuracy.

RALPH also provides a great deal of information, more than is listed in Table 2-1. Computing this information requires time, and thus, has the overall effect of reducing the processing speed of the whole system. Regardless, RALPH was able to operate at up to 25Hz, and provide state output at up to 25Hz. However, system constraints dictated that data be captured at around 15 Hz.

2.3.2 Methodology

Before beginning the study, I needed permission from the Carnegie Mellon Human Use Review Panel (HURP), which oversees any experiments which make use of human subjects. This involved writing a proposal, provided in Appendix B detailing the goals and methodology of the study, along with any potential benefit or harm to the subject. The HURP committee then reviewed this proposal, and made a decision regarding whether to allow the experiment to take place. This study was considered fairly innocuous, and therefore, permission was rapidly granted. Subject recruiting then began, through the use of posters, newsgroup postings, and personal request.

As the goal of the study was to collect data on driving behavior, the experiment involved having the subject drive Navlab 8 from Carnegie Mellon University in Pittsburgh, PA, to Grove City, PA, which is about 50 miles north of Pittsburgh. I was present in Navlab with all subjects. I instructed the subjects to drive as they normally would, to drive safely, and to ignore the presence of the experimenter as much as possible. The only explicit instruction I gave them was to use the turn signal while changing lanes, as the turn signal is instrumented; this allowed automatic marking of data collected during lane changes. While the subjects were driving on the highway, RALPH was active, recording vehicle state information and driver action. The subjects were not told what was being recorded -- only that data on their driving was being stored. After returning to CMU, I provided subjects with an information sheet detailing the purpose of the study, and asked them to fill out a survey regarding their experiences. This survey asks the drivers to rate themselves on their driving performance, and their level of comfort during the study. The results of this survey are presented in Appendix D.

2.3.3 Data Characteristics

This data collection effort took place over seven months, from January to August, 1998. Data was collected on 20 drivers. This section provides statistics on the drivers, along with the data collected, and ends with a discussion of possible confounding effects which may have an impact on the “naturalness” of the observed behavior, and provides a motivation for the follow-on study.

2.3.3.1 Subject Characteristics

Twenty subjects participated in this study. All subjects were CMU students or employees. This constraint was insurance related, as the subjects would be driving a CMU-owned vehicle. Half the subjects were graduate students in the Robotics Institute; the other half were staff members and faculty. However, only two of the subjects were not members (either faculty, student, or staff) of the Robotics Institute. Eight of the 20 drivers were female. The subjects ranged in age from 23 to 43.

2.3.3.2 Confounding Factors

There are a number of confounding factors which could affect the data. This study measured how people drove an unfamiliar vehicle on a highway while in the presence of a potentially intimidating experimenter. The experimental vehicle, Navlab 8, is a mini-van. Only a small percentage of the subjects had any experience driving a mini-van or other large vehicle. Finally, the subjects knew that they were being monitored, and that their performance would be analyzed. This could have lead to more alert, careful driving than might otherwise have occurred.

These effects are difficult to measure in an uncontrolled setting. Properly gauging their contributions to the performance of the subjects would require experiments in which vehicle type and experimenter presence are controlled. Practical considerations did not allow me to perform these experiments. However, it is not unreasonable to suspect that the effects mentioned above did have an influence. Therefore, the data does not necessarily reflect naturalistic driving behavior. These issues are further discussed in Appendix D. Section 2.4 describes a study in which data on natural driving behavior was collected.

While the data may not reflect the subject's natural driving behavior, it does show differences in performance. The controlled nature of the experiments implies that differences in lane keeping ability are due to real differences between drivers, rather than artifacts induced by differences between vehicles or lane tracker calibration.

Another problem with the above data was that, despite the fact that 20 subjects participated, there was not enough data! The training and analysis methods which are explored in Chapter 4 require more data per individual driver than was available. One to two hours of data per driver did not provide enough variability to learn driver response in different situations. The difference in variability by amount of data is shown in Figure 4-3 on page 77. For these reasons, I undertook another study, with the goal of capturing long term natural behavior.

2.4 Naturalistic Data

The goal of second study was to record a large amount (> 5 hours per driver) of natural driving behavior. The motivation was two-fold: To provide more data per driver for my training algorithm, and to reduce or eliminate the confounding factors which may have influenced driver behavior in the previous study. This study resulted in over 70 hours of data on natural driving behavior. This section describes the lane tracker used, the characteristics of the subjects, and the resulting data.

2.4.1 AutoTrak Lane Tracker

AutoTrak was developed by Pomerleau to address some of the issues regarding RALPH which were brought up in Section 2.3.1.2. AutoTrak is significantly smaller and easier to install in a vehicle than RALPH. The interface is a converted radar detector unit. The following sections detail AutoTrak's operation and characteristics. The AutoTrak interface unit and computing box are shown in Figure 2-3, and the state output is given in Table 2-2.

2.4.1.1 Operation

AutoTrak is basically similar to RALPH, in that it uses a camera to look out at the road, and uses lane features to estimate lateral position, lateral velocity, and other state information. Installing AutoTrak is simple. Whereas RALPH is contained in a full-sized PC, the entire AutoTrak computing box is 6" x 6" x 6", and is composed of PC-104 boards. An off the shelf radar detector has been modified, allowing the use of its buttons to control AutoTrak, and the use of its LED panel to display information such as lateral position or menu choices.

Installing AutoTrak involves attaching the radar interface unit to the windshield, directly below the rear-view mirror. The interface connects via a cable to the computing box, which is usually placed on the floor of the rear of the vehicle. The system is capable of self-calibration. This involves driving on a stretch of relatively straight road, where the system is able to determine its camera parameters, such as pitch, over a 30 second calibration period.

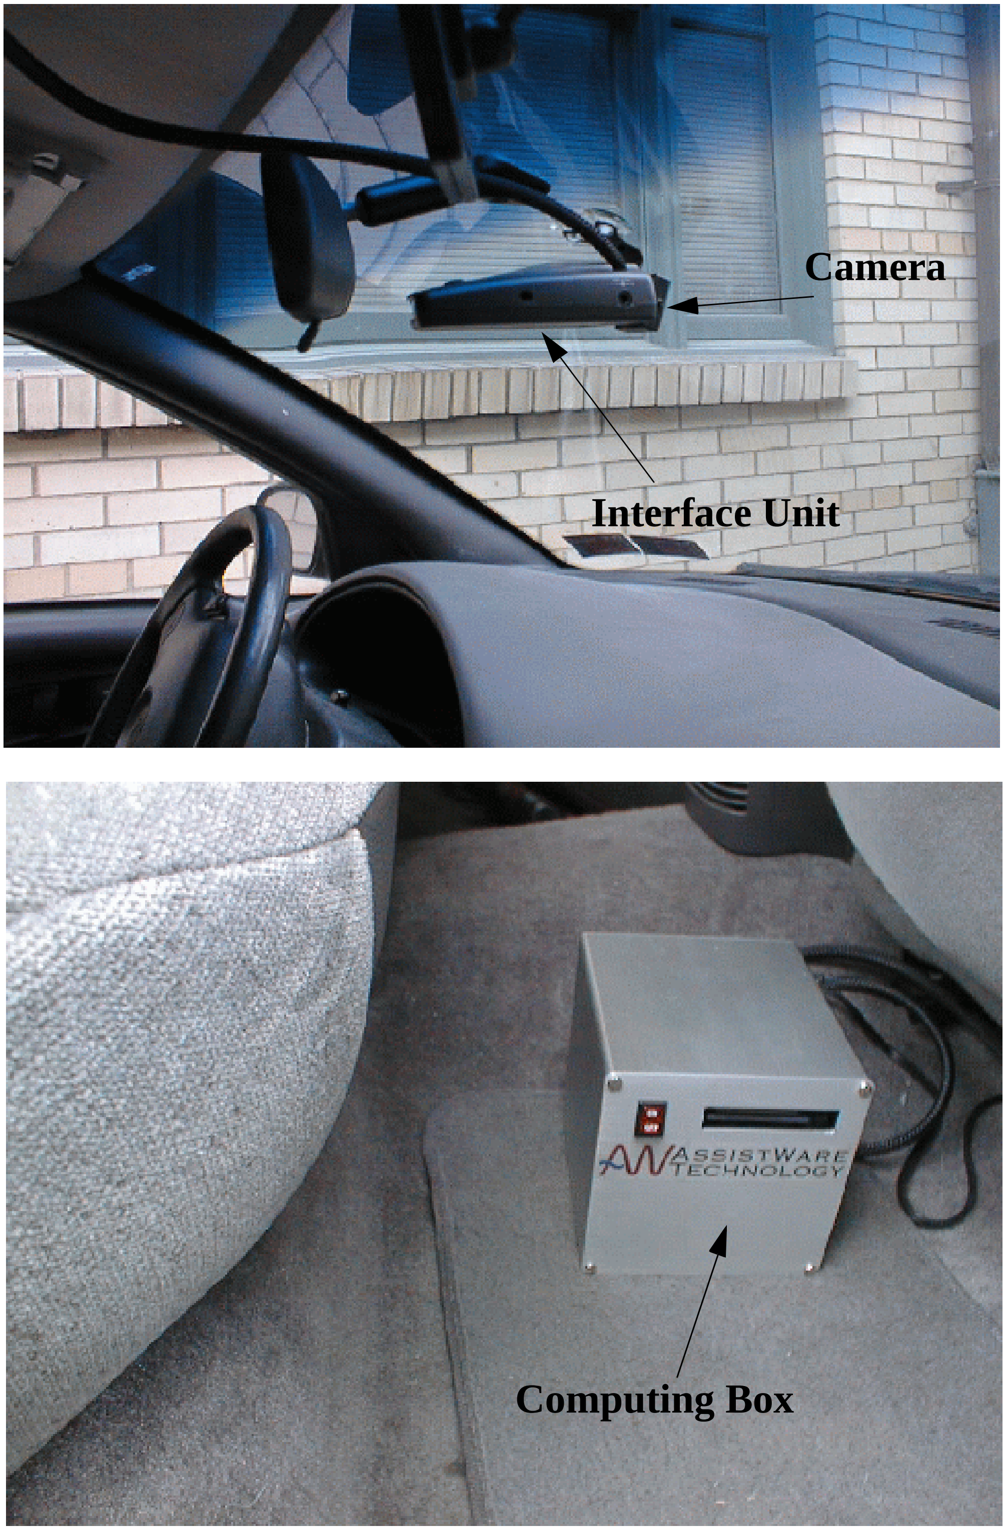

Figure 2-3: AutoTrak interface unit and computing box installed in a subject's vehicle.

State Variable	Range	Units	LSB
Time	0-86,400	seconds	0.01
Lateral Offset	+/-250	centimeters	2.0
Lateral Velocity	+/-250	centimeters/s	2.0
Road Curvature	+/-0.008	1/meters	.000063
Lane Width	260-480	centimeters	1.0
Estimated Lane	left, center, or right lane.		
Offset Confidence	0-100	n/a	1
Turn Signal	left or right turn signal active		
Velocity	0-50	meters/s	0.2
Heading	0-360	degrees	0.1
GPS Latitude	0-90	degrees	.000001
GPS Longitude	+/-180	degrees	.000001

TABLE 2-2 AutoTrak State Output

The system also continuously refines its calibration during normal operation. This ease of installation allowed me to install the system in subjects' personal vehicles, as is described in Section 2.4.2.

2.4.1.2 Algorithm Characteristics

AutoTrak trades off curvature accuracy for lateral position accuracy. While this trade-off prevents AutoTrak from controlling an autonomous vehicle, the road curvature estimates provided are more than adequate for segmenting roads into curves of varying degree. The high cycle rate allows data to be collected at up to 60 Hz, although in this case, data was collected at only 30 Hz due to storage constraints.

2.4.2 Methodology

As in the Navlab 8 data study, I had to obtain permission from the university human use review panel before beginning the data collection. Permission was again rapidly granted, as in this case, the installation of AutoTrak would provide no distraction or risk to the driver, because the system would be operating in pure data collection mode. For safety, as well as methodological reasons, all audible and graphical output was turned off. In normal operation, AutoTrak will sound an audible alarm when drivers cross over a solid lane boundary. As of yet, there have been no studies done on the effect such a warning would have on drivers. It is possible, although perhaps unlikely, that such an active intervention could startle them, causing them to react inappropriately. In addition, this data collection effort was conducted to collect examples of *natural* driving behavior. Having an active warning and graphical depiction of lane position could affect the actions of the driver. Anecdotally, there is evidence that having the graphical lane position indicator active causes drivers to center themselves more accurately when driving an unfamiliar vehicle.

After I received permission to begin, I recruited subjects via e-mail and newsgroup postings. The majority of the collection occurred over the Thanksgiving and Christmas breaks, as many people go on long trips during these holidays. The study required the subjects to bring their car to CMU, where I installed the AutoTrak system. After installation, each subject and I went for a short drive to calibrate the system and familiarize the driver with AutoTrak. During this period, the audible and graphical outputs of AutoTrak were left on, to verify calibration. After calibration was completed, the audible and graphical outputs were turned off. The system was put into data collection mode, in which it would record data whenever the road conditions allowed it to generate a valid estimate of vehicle state. Subject involvement with the system included only inserting the power cord into the cigarette lighter before a drive. Many of the subject vehicles had cigarette lighters which deactivated when the car was off, meaning that the subject did not have to do anything.

When the subjects returned to CMU, AutoTrak was removed, and the subjects were each paid $50 for their involvement in the study. The subjects were then told (if they indicated interest) what had been recorded, and the purpose of the study.

2.4.3 Data Characteristics

This data collection effort took approximately 3 months to complete, from November 1998 to January, 1999. The study resulted in data being collected on nine drivers, resulting in about 70 hours of data collected. This section provides statistics on the subjects and ends with a discussion of possible confounding effects which may have an impact on the integrity of the data.

2.4.3.1 Subject Characteristics

Nine subjects volunteered for this study. As in the previous experiment, insurance concerns dictated that they all be members of the CMU community. Three of the subjects were female. The subjects ranged in age from 25 to 43. Six of the subjects were graduate students at CMU.

2.4.3.2 Confounding Factors

The main confounding factor during this study was a general AutoTrak installation error which was noticed after the data collection was complete. This error was discovered after an attempt to measure the lateral position and lane width estimate accuracy of AutoTrak.

While the curvature accuracy is difficult to measure, it is possible to measure lateral position accuracy. To do this, I conducted an experiment with AutoTrak in the parking lot of the Pittsburgh Zoo. I set up and measured artificial lane markings. The lane markings were set up by extending a line from the center of the vehicle, and placing markings parallel to that line, at a distance of 1.8m, for a total lane width of 3.6m. (See Figure 2-4, which depicts the

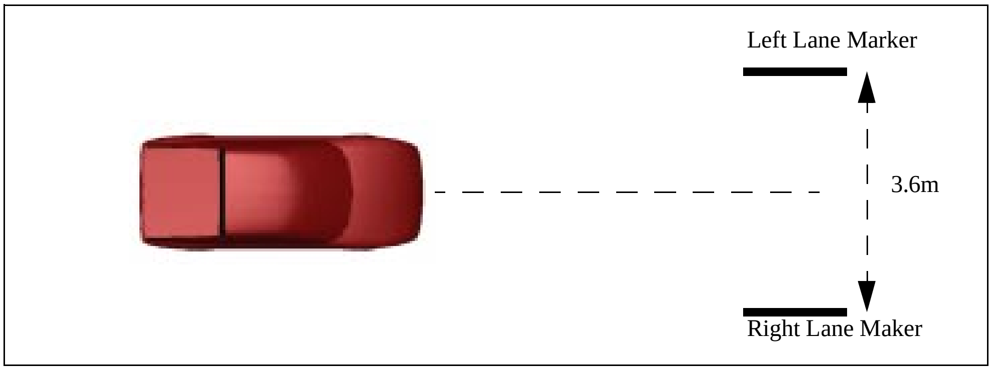

Figure 2-4: Top-down View of AutoTrak Precision Experiment

setup used.)

I tested lateral position accuracy by manually moving the lane markers 5cm at a time, taking lateral position readings at each step. This was done over the range [-1.2m, 1.3m] (where a negative offset is left of center, and positive offset is right of center). Similarly, I tested AutoTrak’s ability to measure lane width by varying the distance between the artificial lane markings over the range [2.4m, 4.16m]. The results are shown in Figure 2-5.

Figure 2-5 shows the estimates and estimate errors for lateral position and lane width. The top two graphs show the AutoTrak lane position and lane width estimates overlaid on the actual lane position and lane width. The bottom two graphs show the errors in lane position and lane width estimates. The overall error is quite low, and nearly linear. In the range [-50, 50] cm, the error varies from 0 to -8 cm. The graph shows that AutoTrak tends to over-estimate the lateral position, making a vehicle seem further away from the lane center than it really is. While this could be a source of false alarms, it is not. One reason is that the lane width estimate is also over-estimated, between 36 and 56 cm.

These state variables were being over-estimated because of a false assumption that AutoTrak makes regarding camera height. To allow for online calibration of camera pitch angle, either lane width or camera height must be known. AutoTrak assumes that, when the system is installed, lane width can vary but camera height will remain stable. Therefore, AutoTrak uses numerous vehicle classes, such as sedan, mini-van, SUV, and truck. Each of

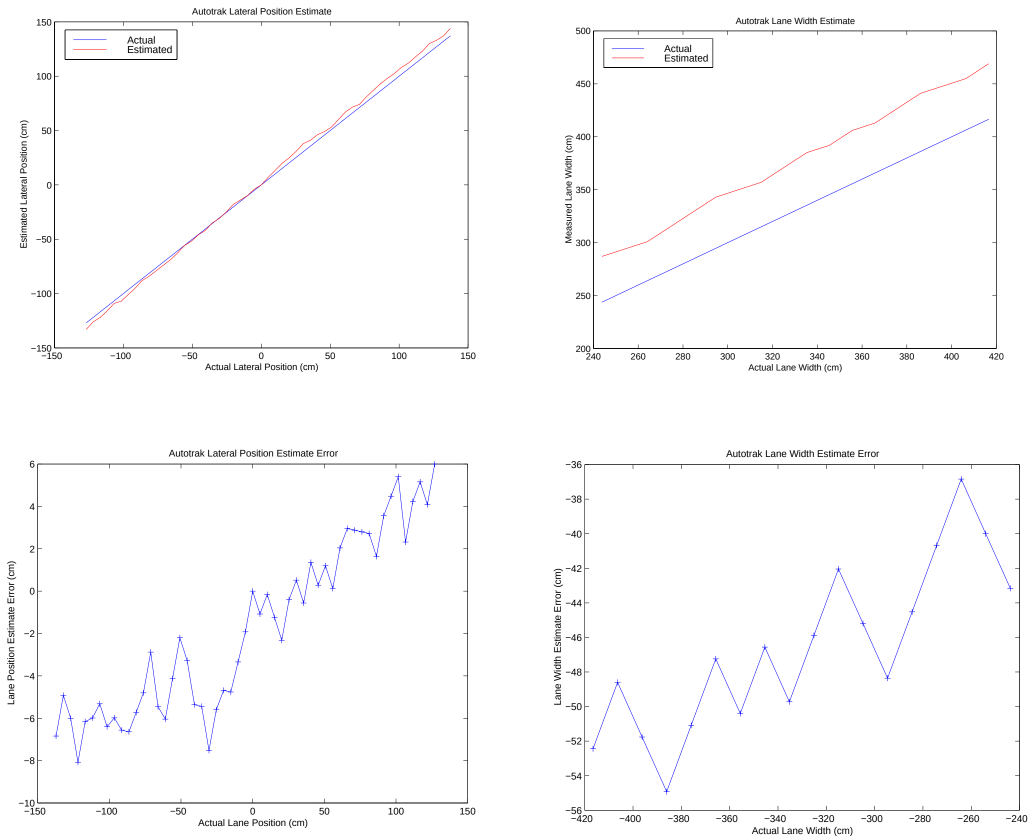

Figure 2-5: AutoTrak Lane Position and Lane Width Estimates and Error.

these vehicle types has an associated camera height. This means that camera height did not have to be measured at each installation, based on the assumption that each sedan (for instance) has its rear-view mirror at the same height. During installation, AutoTrak is told which type of vehicle the system is being installed in. The problem occurred in AutoTrak's assumption of rear-view mirror height for sedans, which was over-estimated. Therefore, the camera was really installed lower than the system assumed. This caused the image to appear wider, and caused over-estimates of lane width and lateral position.

Over 70 hours of driver data was collected before this problem was noted. Fortunately, it does not have an impact on comparing lane departure warning algorithms within driver, as I only care about relative performance in that case. It does, however, create an issue in comparing performance between drivers, as the height of the camera did change slightly between installations in different vehicles, causing a small error in lane position and lane width estimates. For this reason, the data cannot be corrected based on the above error measurements. However, as the overall lane position error is relatively low, and in general agreement with the increase in lane width, the calculations of warning system performance and analyses which are presented in the next chapters are not adversely affected. The only corrections which could be applied would just be a linear scaling of lane position and lane width, and therefore would not change the gross behavior of any of the warning algorithms evaluated in this thesis.

Confounding factors which existed in the Navlab 8 study were not an issue in this experiment. The subjects were driving their own vehicle, and no experimenter was present. More importantly, the subjects were engaged in long drives, often with family or other passengers within the car. The presence of other passengers can help drivers forget that they are being monitored, and should encourage natural driving.

2.5 Summary

I begin this chapter with a motivation for the use of real world data, and then describe the details of two driver studies which were conducted using lane tracking systems. The first study had a more controlled environment, as it occurred in a single vehicle (NavLab 8), around the same time of day, and in the presence of an experimenter. The second, naturalistic study, was more free-form. It involved installing an AutoTrak data collection system in the vehicle of a volunteer who was going on a extended (greater than 500 miles) trip. The experimental methodology, resulting data, and confounding factors of both studies are discussed. This chapter lays the foundation for the presentation and discussion of warning algorithm performance, which I present in the next chapter.

CHAPTER 3 Lane Departure Warning

3.1 Introduction

In this chapter, I present a model for predicting dangerous lane departures and explore whether any differences in driver behavior can be exploited to improve the performance of a lane departure warning system based on such a model. The purpose of this is to predict when drivers are about to do something dangerous, and warn them in time so that they can take corrective action. I am not interested in models of how people drive down the middle of the lane. There has been a great deal of previous work in that area. Rather, I am concerned with models of lane departures. In particular, I focus on lane changes, as they are a form of lane departure which is relatively easy to capture.

A warning system should give as much warning time as possible, while triggering few, if any, nuisance alarms (as defined in Section 1.2). The lane departure warning algorithm I use has two parts: a vehicle state predictor, and an alarm decision model, which are defined in Section 1.2. This algorithm has five parameters. This is a consequence of driver variation during lane changes. If all drivers changed lanes with the same lateral velocity in all situations, it would be possible to produce a model with just one parameter. This parameter could be adjusted to increase the amount of warning time, with a corresponding increase in the number of nuisance alarms. This parameter could then be tied to a knob marked “sensitivity,” and left to the user to adjust to a level at which the number of nuisance alarms was acceptable. However, this is not the case. Drivers vary their lateral velocity when changing lanes, as shown in Figure 4-3 on page 77. The marked off lane change areas in this figure show a spread of lateral velocities. There is a further issue in that drivers do not drive in a straight line, i.e., their lateral velocity during the entire lane change maneuver is not constant. Godthelp [24] models lane changes as a sinusoidal trajectory. However, this is not completely accurate. True lane change behavior lies somewhere in between a linear and sinusoidal model, and is not consistent within drivers. Figure 3-12 illustrates a right lane change and the associated lateral velocity. The lateral velocity increases for about 1 second, then stays constant for about 2 seconds before the lane change event. This variability makes lane changes difficult to analytically model, and requires looking at real world data. As I show, differences between drivers, within drivers over time, and over road geometry require more intelligent modelling of the driver's behavior. These differences do not follow a single mathematical model. Due to this, I explore training some of the parameters of a more sophisticated warning algorithm for individual drivers.

I begin with a discussion of previous work based on two current methods: Roadside Rumble Strips and Time to Lane Crossing (TLC). Following that, I present a description of the Future Offset Distance (FOD) algorithm, which is an extension of TLC. I then show the results on real world data of roadside rumble strips and of a standard TLC algorithm, to demonstrate the problems in warning time and excessive nuisance alarm rate the two methods have, respectively.

After that, I present results on training an FOD based warning algorithm tailored to an individual driver, and compare those with two other models: 1) a generic FOD model which is tailored to a set of drivers, and 2) the FOD model which the AutoTrak lane departure warning system uses. Then I present model extensions which can account for short-term changes in driver behavior due to variations in road geometry and variations over time; I then do a parameter analysis of the parameters involved. These model extensions improve warning system performance by lowering the number of nuisance alarms from 163 to 90, in the simplest case, while leaving the warning time mostly constant. Finally, I describe the behavior of the

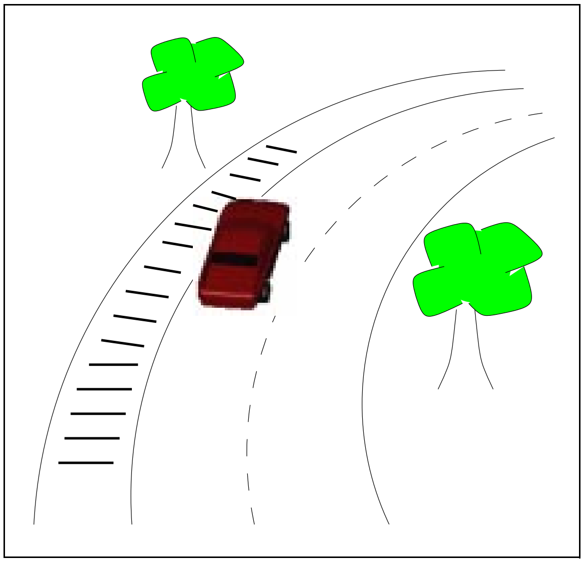

Figure 3-1: Roadside Rumble Strip System

individually trained models, and discuss some of the insights that these experiments have provided regarding driver behavior. All the experiments and analyses in this section were done using real world data collected in the naturalistic study described in Section 2.4.

3.2 Previous Warning Algorithms

There is a great deal of previous work on driver models, reaching back to the 1950s. Most of these models take some combination of vehicle state and road geometry as input, and output the driver's expected steering command. This previous work, along with an overview of warning systems, is presented in the introduction, in Section 1.3. In this section, I focus on two specific warning methodologies: roadside rumble strips and TLC. This is in preparation for experiments in the following section, which will demonstrate weaknesses of both methods.

3.2.1 Roadside Rumble Strips

This is a warning system which relies on grooves cut in the shoulder 15-45cm. beyond the physical lane boundary. Figure 3-1 illustrates this. Drivers hear a low frequency sound and their vehicle vibrates when a tire rolls over the grooves. roadside rumble strips has been

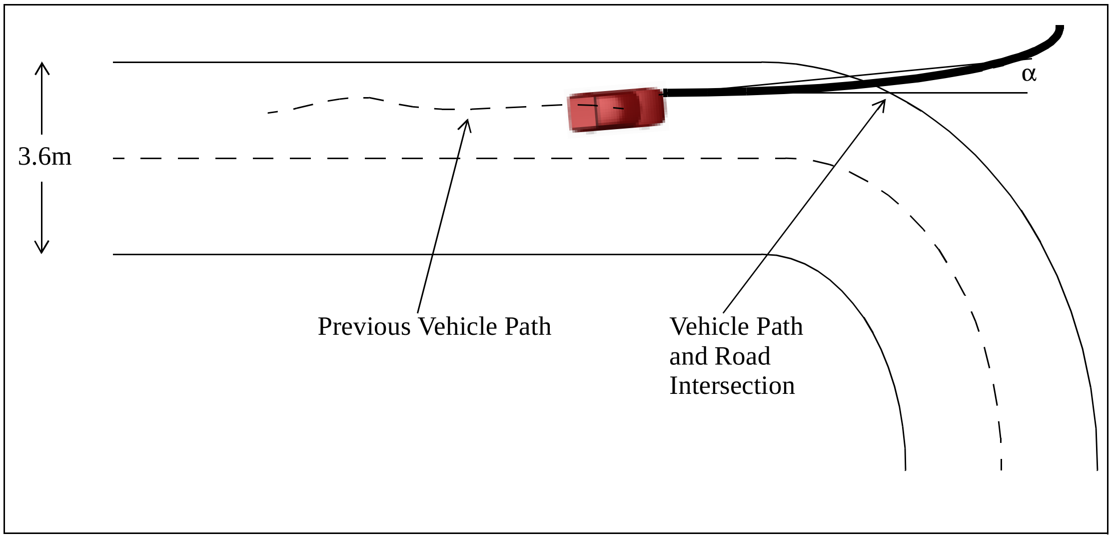

Figure 3-2: Depiction of the Time To Lane Crossing (TLC) algorithm. The TLC is a measure of the time it would take for the first tire of a vehicle to cross a lane boundary. The prediction of the vehicle's path assumes that there will be no change in steering over the length of the prediction, and can be computed using the vehicle's velocity, relative yaw, α, and upcoming curvature.

shown to reduce the number of run off road accidents by 70% [75]. There are, however, problems with roadside rumble strips. The first is that it is infrastructure-based. It is not present on all highways, and adding it can be expensive. The second problem is that the amount of warning the driver gets can be low because the alarm doesn't trigger unless the driver actually *is* in danger. Finally, the number of nuisance alarms is high for drivers who tend to weave a lot. This is especially problematic for truck drivers, who tend to weave quite a bit. Results on evaluating roadside rumble strips on real driver data are presented in Section 3.5.2.

3.2.2 Time to Lane Crossing (TLC)

The Time to Lane Crossing metric, which was first proposed by Godthelp [24], is a measure of the time remaining before a vehicle on a given trajectory will depart the road. Figure 3-2 illustrates this. The vehicle is entering a right curve, but is not yet steering to follow the curve perfectly. At this instant, the TLC is the amount of time the vehicle has before the

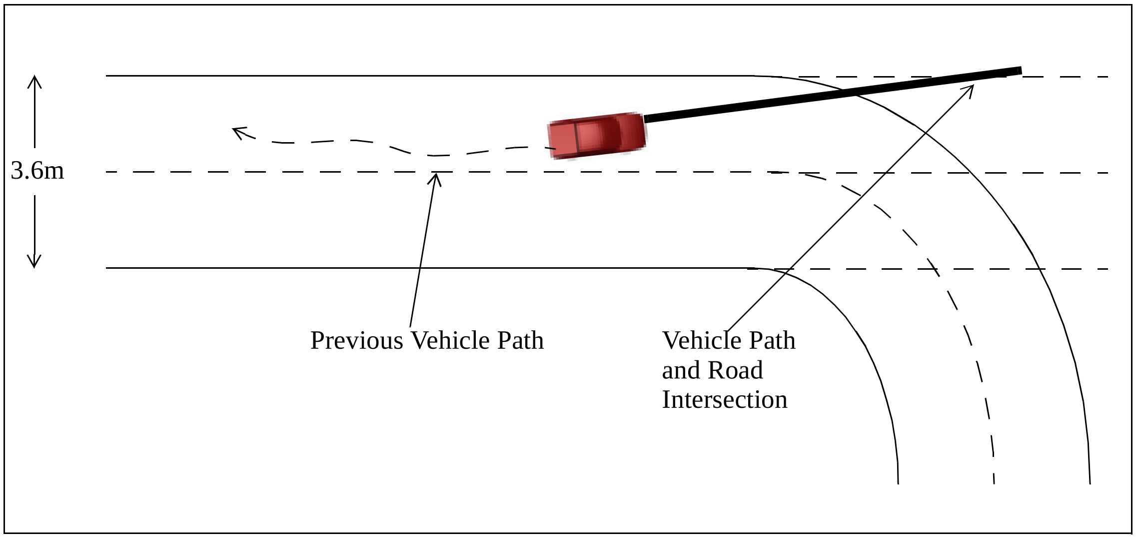

Figure 3-3: Simplified TLC model. The vehicle is assumed to be traveling in a straight line based on its forward velocity and heading, and the road is assumed to be locally straight. Therefore, TLC is proportional to the length of the vehicle path to the lane boundary. Although the consequence of ignoring road curvature is exaggerated in this illustration, the true effect is usually minimal due to the relatively gentle curves on most highways, and small vehicle heading angles (generally 0-1.5 degrees).

front-left tire touches the left lane boundary. TLC generally assumes that the vehicle will not change its steering radius over the length of the predicted path, and is calculated (by [41]) by predicting the vehicle's future path using discrete timesteps.

The accuracy of TLC can be improved through the use of more complicated vehicle models, for increased accuracy in path prediction, as shown by Lin [41] and others at the University of Michigan. Conversely, the vehicle and road models can be simplified, to use only information which is available from a vision system, as was done by Pomerleau. Figure 3-3 illustrates this approach.

Rather than iteratively predicting the vehicle's path and comparing it with the lane boundary, the simplified TLC is computed in one step. This is done by assuming that the vehicle has a constant lateral velocity, determined by state measurements with the road. The other assumption is that the road is locally straight. The yaw, combined with forward velocity and lateral position, can be used to calculate the time before the first wheel crosses the lane boundary. In the situation depicted in Figure 3-3, we see that these assumptions lead to an over-estimation of TLC, which could potentially delay the triggering of a true alarm. However, in practice, these errors are quite small due to the relatively shallow nature of most highway curves and low relative yaw displayed by vehicles.

In general, TLC provides more warning time than roadside rumble strips, because warnings are triggered when a driver is *predicted* to be in danger. However, this prediction can be wrong, and therefore, the number of nuisance alarms is generally much higher than with roadside rumble strips, as I show in Section 3.5.3. In the next section, I describe the Future Offset Distance (FOD) algorithm, which provides much of the increased warning time of TLC, while maintaining nuisance alarm rates similar to roadside rumble strips.

3.3 Future Offset Distance

This section describes the Future Offset Distance algorithm which I use as the basis for my lane departure warning experiments. This model allows the driver to drift past the physical lane boundary before being alerted to possible danger. This is a more realistic model of driver behavior, as certain drivers tend to drift beyond the lane boundary during normal driving. Allowing this drifting to occur reduces the number of nuisance alarms. Although the roadside rumble strips system recognizes this by placing the rumble strips beyond the lane boundary, roadside rumble strips triggers when the driver is actually *in* danger, defined as being 15-45 cm. beyond the lane boundary. As these rumble strips are normally only used on wide shoulders, the nature of the danger is debatable. This reduces the amount of reaction time available to the driver.

Unlike roadside rumble strips, conventional TLC triggers an alarm whenever the vehicle is predicted to exceed the physical lane boundary (Pilutti's Virtual Rumble Strips [58] are an exception). This allows for a prediction of danger, which allows the driver more reaction time than roadside rumble strips. However, by triggering a warning when the vehicle is predicted to cross the physical lane boundary, TLC can trigger many nuisance alarms. FOD is similar to TLC in that it predicts when the driver is in danger of leaving the road, and triggers an alarm when the time to lane crossing is below a certain threshold. However, FOD extends roadside rumble strips's idea of allowing the driver to drift beyond the lane boundary by adding a virtual lane boundary. If all drivers exceeded the lane boundary by the same amount, then the width of the virtual lane boundary could be fixed, and all would be fine. However, different drivers have different magnitudes of swerving, which means that the virtual lane boundary needs to be a parameter. The advantage of FOD over roadside rumble strips is that the virtual lane boundary can vary, and can be adjusted to suit an individual driver. If a particular driver tends to always stay within the lane, then the virtual lane boundary can be set appropriately. More warning time could be provided to such a driver, as opposed to a driver who normally drifts 10-20 cm. out of the lane. The FOD algorithm provides a framework in which to explore and exploit such differences in driver behavior.

3.3.1 Parameter Description

There are two parameters in the FOD algorithm: lookahead time, *T*, and virtual lane boundary, *V*. The FOD lookahead time is how far in the future the system is willing to predict future vehicle state. The virtual lane boundary is a distance beyond the physical lane boundary which the driver is allowed to occupy. Figure 3-4 illustrates this. A warning is not triggered when the vehicle is predicted to exceed the lane boundary. Rather, it is triggered when the vehicle is predicted to exceed the *virtual lane boundary*, at time *T* from now, which is usually beyond the physical lane boundary. If the following relation is true, then an alarm is triggered:

$$L_p' > V \tag{3-1}$$

Where L_p' is the predicted vehicle lateral position, and V is the virtual lane boundary. There are many ways to compute the predicted lateral position, L_p'. I use a 1st order kinematic approach, where:

$$L_p' = L_p + TL_v \tag{3-2}$$

Where L_p is the current distance from the virtual lane boundary, L_v is the lateral velocity, and T is the lookahead time. Other models are possible, such as 2nd order models which use lateral acceleration information, or data-centric approaches such as memory based learning which use a distribution of actual future lateral positions calculated using training data.

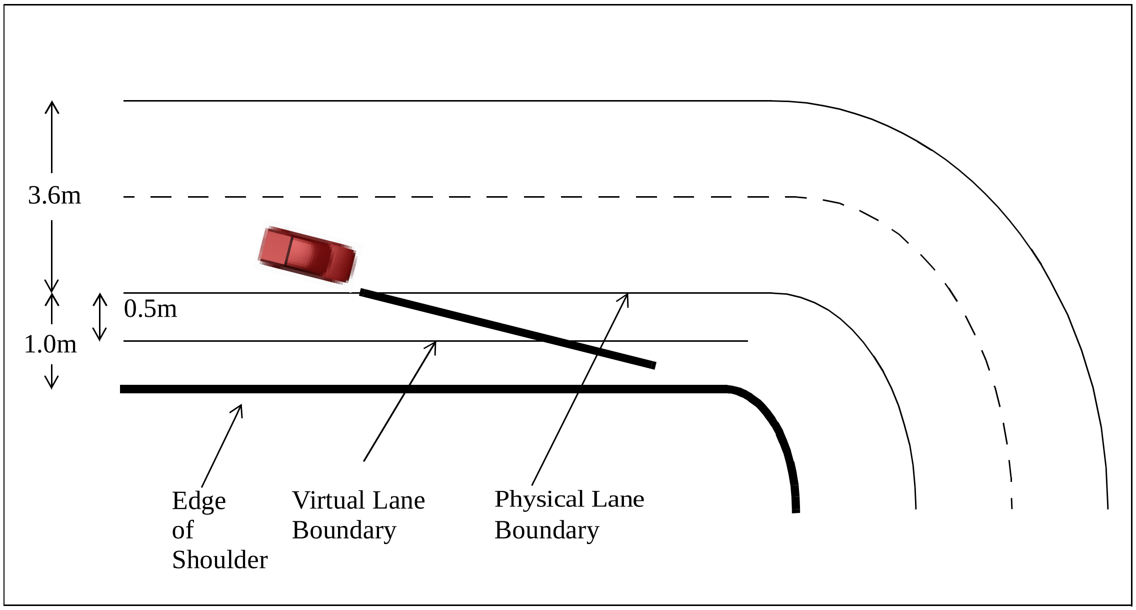

Figure 3-4: Future Offset Distance (FOD) Algorithm. An alarm triggers when the vehicle is predicted to cross a virtual lane boundary.

3.3.2 Performance

Mathematically, there is a set of FOD parameters which generate identical behavior, assuming a constant lateral velocity. Figure 3-5 shows a case which illustrates this. If we assume that the vehicle has a constant lateral velocity L_V, and we want to trigger an alarm when the right tire touches the lane boundary, then there are an infinite number of FOD lookahead time and virtual lane boundary parameters which will produce the desired behavior. Intuitively, if we increase the virtual lane boundary, then to trigger at the same point, the lookahead time has to be increased by a proportion of the lateral velocity:

$$\frac{V - x}{L_v} = T \tag{3-3}$$

where V is the desired virtual lane boundary, x is the desired trigger point (in meters, relative to the lane edge), L_V is the lateral velocity, and T is the lookahead time. This relationship shows that, for any desired trigger point and virtual lane boundary, a lookahead time can be found which satisfies the conditions.

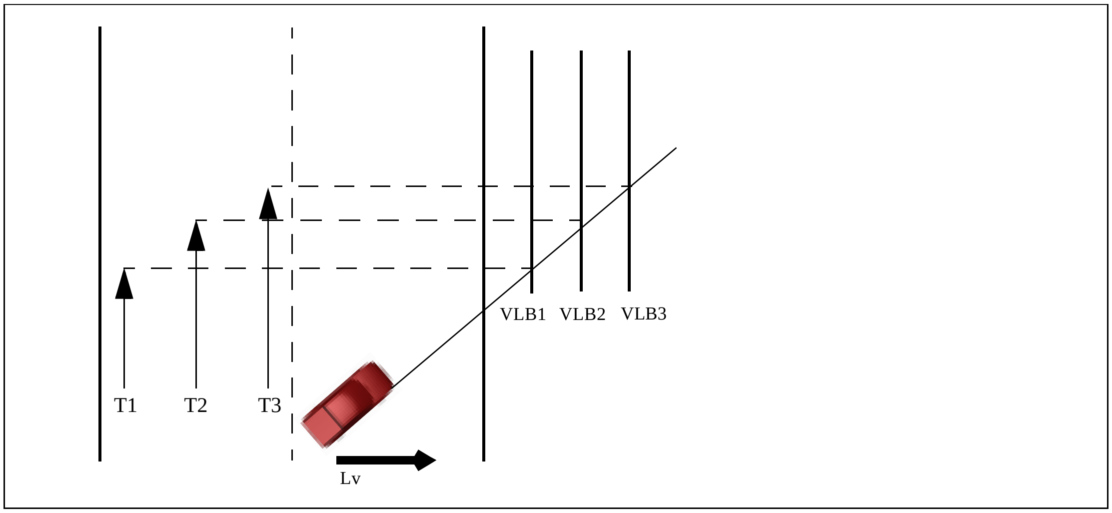

Figure 3-5: Depiction of FOD parameter equivalence. For a particular lateral velocity, there is an infinite set of VLB and Lookahead Time (*T*) parameters which trigger an alarm at the same location. The FOD parameters (T1, VLB1), (T2, VLB2), and (T3, VLB3) will all trigger at the same location.

While Equation (3-3) shows that there is an infinite set of equivalent FOD parameters, in practice, not all parameters provide equal performance. As we will see later, the driver's lane keeping ability plays a role in parameter selection. If the driver tends to weave a lot, then a small virtual lane boundary value can trigger nuisance alarms. In these cases, increasing the virtual lane boundary and the corresponding lookahead time can provide the same amount of warning, while lowering the number of nuisance alarms. In the next section, I present experiments which explore this idea.

3.4 Experimental Design

The purpose of the experiments in this section is twofold: 1) To evaluate previous methods using real world data, and 2) to determine whether differences between drivers can be exploited to improve the performance of a future offset distance based warning algorithm, by tuning the FOD parameters for a specific driver. The motivation for doing this is to try to reduce the number of nuisance alarms, while still maintaining an adequate warning time. Current systems, such as AutoTrak, assume that a user adjustable sensitivity setting provides adequate performance for all drivers. There are two questions inherent in this assumption: 1) If one set of parameters is adequate for all drivers, what are those parameters? and 2) Does one set of parameters adequately represent all drivers? In particular, I look at the effectiveness of a warning system based on the AutoTrak FOD parameters, vs. a warning system which is trained on a “generic” driver, which is a combination of all drivers (to examine question 1) vs. an FOD model trained for an individual driver (for question 2). The results will show that in certain cases training does provide better performance.

3.4.1 Terminology

In the rest of this chapter, I will be using the following notation:

- *D(n)* - Driver “n”.
- *Dgen(n)* - This is a “generic” driver, which is a time-weighted combination of all the driver data *except* that of driver “n”.
- *T* - FOD (or TLC) Lookahead Time.
- *V* - FOD Virtual Lane Boundary.
- *FOD(n, w)* - Trained FOD parameters with the lowest nuisance alarm rate given the desired warning onset time *w*.
- *FOD(T, V, n)* - Apply the FOD model with parameters “T” and “V” to driver “n”
- *wot(FOD(T,V))* - The WOT resulting from applying parameters “T” and “V”.
- *nar(FOD(T,V))* - The NAR resulting from applying parameters “T” and “V”.

3.4.2 The Data

The naturalistic driver study described in Section 2.4 involved nine subjects, and collected about 70 hours of data. However, not all of the data was useful for these experiments. Of the nine subjects, three had data which was unusable due to an error which resulted in 10-30 second dropouts every 3 minutes or so. An additional driver had long stretches where the overall confidence of the AutoTrak system was low. This subject indicated that most of the data was collected on either snowy or rural roads, which could explain the extended periods of low confidence. Finally, the data for Driver 1 contained an unknown amount of data from a different driver. This is because the subject allowed a passenger to drive for a short while, without turning off the data collection system. The subject has indicated that he did the majority of the driving, and that only a small portion was done by the passenger. Therefore, I use the data, with the caveat that the results on this driver may be colored. The total amount of data

Driver	Amount of Data (hours)	Number of Lane changes
Driver 1	5.22	170
Driver 5	2.76	67
Driver 7	1.44	55
Driver 8	2.54	93
Driver 9	6.54	219
Total:	**18.50**	**604**

TABLE 3-1 : Details of the data used in this experiment

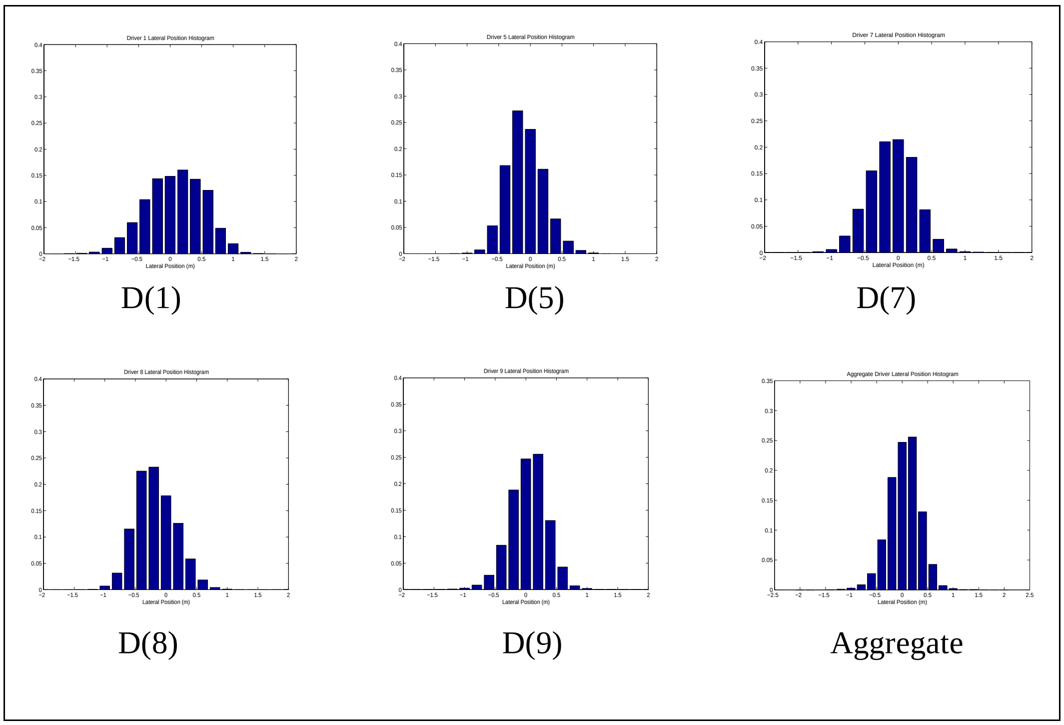

Figure 3-6: Lateral Position Histograms of Test Subjects.

used is listed in TABLE 3-1. Of the nearly 70 hours of data, only 18.5 were usable. The number of lane changes is important, because, as I explain below, they will be used as surrogate lane departure events.

Figure 3-6 shows a histogram of lateral position for the five test drivers, along with a composite distribution of all the drivers. Figure 3-7 shows the same information for lateral velocity. Note that most of the histograms look alike. This is not surprising, as there is proba-

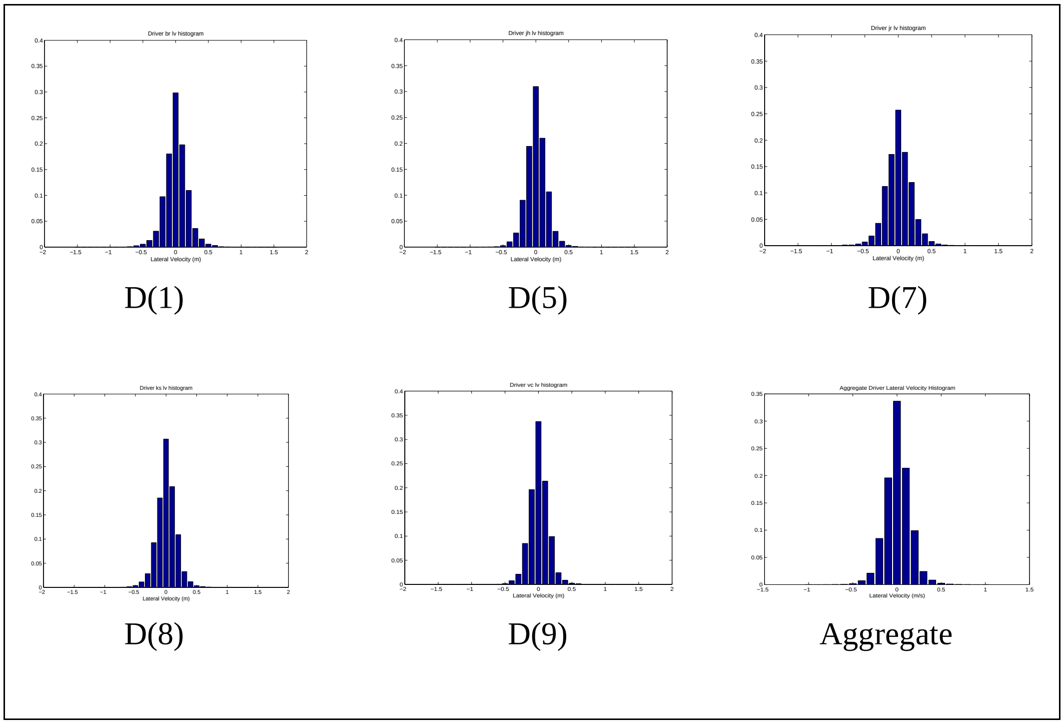

Figure 3-7: Lateral Velocity Histograms of Test Subjects.

Driver	Lateral Position Mean (m)	Lateral Position Standard Deviation (m)
1	0.08	0.45
5	-0.08	0.29
7	-0.09	0.34
8	-0.17	0.33
9	0.04	0.30

TABLE 3-2 : Lateral Position Means and Standard Deviations

bly fixed range of comfortable lateral velocities when driving. TABLE 3-2 shows the means and standard deviations of the lateral position for the five drivers. The lateral position data from three seconds before and three seconds after each lane change was discarded when calculating these statistics. A negative lateral position is left of center, and positive is right of center. The spread between the left-most and right-most mean lateral position is small -- about

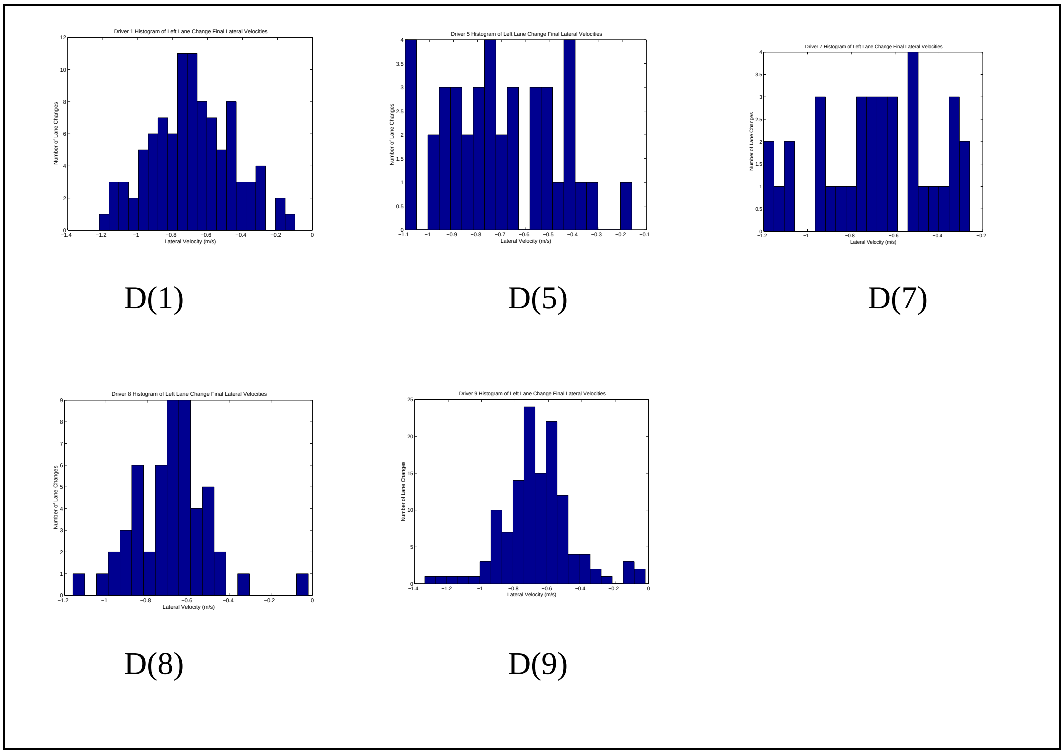

Figure 3-8: Lateral Velocity at Lane Change event for Left Lane Changes

0.25m. Most of the standard deviations are similar as well, except for D(1), whose standard deviation is higher than that of the other drivers. We will see later that the difference in standard deviation will have an impact on warning system performance.

The standard deviations reported here are somewhat larger than statistics which have been seen in other work. [57] has a summary of these statistics, which show lane keeping standard deviations ranging from 0.14m to 0.24m for different types of vehicles (cars, minivans, and trucks) under different circumstances (test track or freeway driving). One possible reason for the discrepancy is that the individual driver statistics I present were computed on larger amounts of data than previous work, which was collected in changing environments and over discontinuous periods of time, allowing for individual drivers to display varying behavior.

Figure 3-8] and Figure 3-9 show histograms of lateral velocities at the point where AutoTrak detects a lane change. Note that the majority of the lane changes like in the 0.5m/s to 1.0m/s range. This provides evidence that drivers do not always change lanes with exactly the same trajectory.

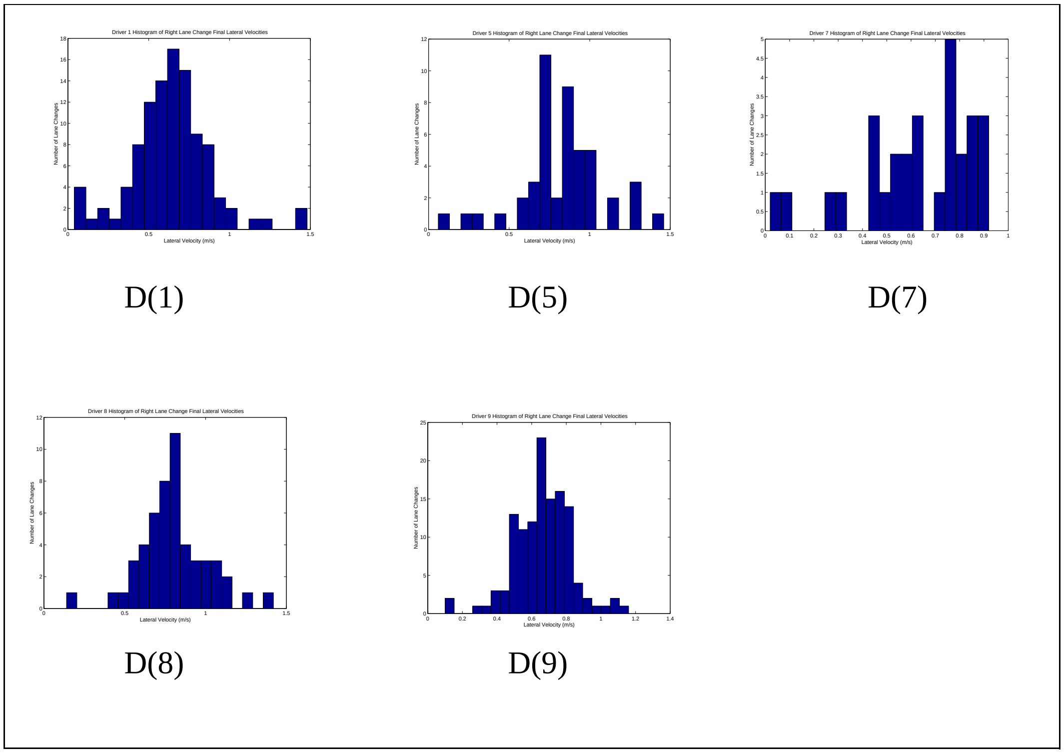

Figure 3-9: Lateral Velocity at Lane Change event for Right Lane Changes

3.4.3 Evaluation

This section describes the methodology of the experiments, which involves testing roadside rumble strips and TLC performance, and finding the optimal FOD parameters for a particular driver and comparing it against a generic model and the hand-tuned AutoTrak model. The warning algorithm is evaluated using two metrics: Warning Onset Time (WOT), and Nuisance Alarm Rate (NAR). WOT is the time between the alarm trigger and the first tire touching the edge of a pre-set virtual shoulder, and is calculated assuming lane changes are true alarms. The nuisance alarm rate measures how many nuisance alarms occur per hour.

I chose to use lane changes as surrogate lane departures because true lane departures which involve an accident are very rare. Even if the data contained a true road departure, it would be difficult to evaluate system performance using just one example of a road departure. Therefore, I am using lane changes, which are similar to certain types of road departures, primarily those due to unintended steering wheel motion. Lane changes are not similar to other types of road departures -- particularly those in which the driver slowly drifts off the road. A lane change generally involves lateral velocities of 0.5-1.0m/s, as shown in the memory table in Figure 4-2 on page 76, where the lane change areas lie in this lateral velocity range. Slow drifts, however, can have lateral velocities as low as 0.2m/s.

3.4.3.1 Alarm Generation and Suppression

When the vehicle is predicted to be beyond the virtual lane boundary, an alarm is triggered. The location of this alarm is stored for later analysis. In the case of a non-lane change alarm, it is possible that the vehicle will be in an alarm generating state for an extended period of time. Rather than trigger a separate alarm for each timestep, I use an alarm suppression heuristic which generates one alarm per excursion event. By excursion event, I mean a contiguous period of time in which the vehicle is in an alarm state. For example, if the driver is entering a construction zone, and is forced to one side of the lane by traffic cones, it would be inappropriate to continuously sound an alarm. Rather, it is preferable to sound an alarm once, and then not trigger it until the driver has entered a quantitatively different excursion state. It would be best if the system were to detect that the vehicle is in a construction zone and not trigger at all; however, that is beyond the scope of the work described in this thesis. An approximation to this is discussed in Section 3.6.2. When an alarm is triggered, all future alarms are suppressed until the driver has not been in an alarm state for the previous six seconds. This prevents clustering of multiple alarms, and is shown in Figure 3-10.

3.4.3.2 Alarm Classification

Not all alarms are true alarms. Drivers who weave a lot trigger nuisance alarms, and detecting and classifying these events is important. Using lane changes as surrogate lane departures simplifies this task. It is easy to detect lane changes in AutoTrak data, as there is a signal which indicates when AutoTrak thinks a lane change has occurred. This is based on a combination of lateral position and lateral velocity and the acquisition of new features in the adjacent lane. The lateral position data during this time also jumps, as shown in Figure 3-12. Combining the Autotrak lane change event with discontinuity detection allows for accurate tagging of actual lane changes.

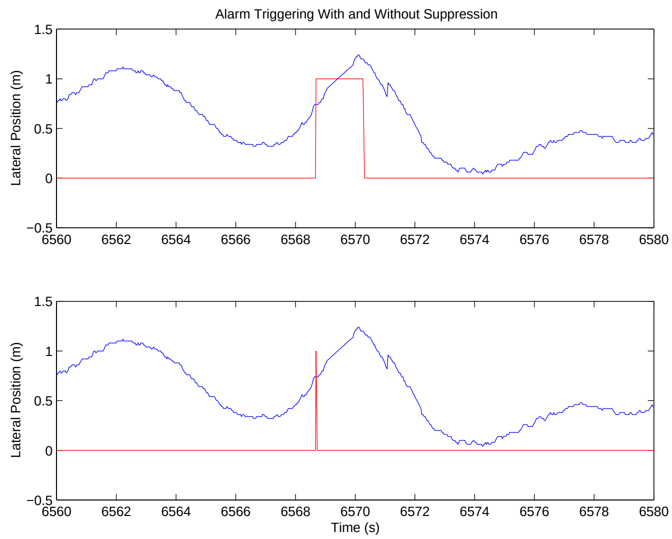

Figure 3-10: Demonstration of Alarm Suppression. The top plot shows a stretch of time where an alarm would trigger, indicated by the lighter square wave. The bottom plot shows how the alarm suppression heuristic would handle this case.

The classification then operates as follows: For a given driver D(n), the FOD algorithm is run over all the data for that driver. This results in a set of alarm triggers, *S*. For each alarm trigger, $a \in S$, the neighborhood of *a* is searched. If a lane change is found, then *a* is marked as a true alarm. In all other cases, the alarm is marked as a nuisance alarm. The definition of nuisance alarms in Section 1.2 states that nuisance alarms are alarms which the driver does not believe were necessary. That is impossible to judge when working with offline data, so I use a relatively harsh definition of nuisance alarms: namely, anything which is not a lane change, is a nuisance alarm. It is likely that in an online system, drivers may feel that some of the alarms which are classified as nuisance alarms, particularly those which involve a large (greater than 1 ft.) excursion, are justified, and should be classified as true alarms. I do not make such a judgement because there have been no studies to determine what drivers consider nuisance and true alarms. Therefore, the nuisance alarm rates which are reported here may be higher than actual nuisance alarm rates.

3.4.3.3 Warning Onset Time Calculation

The Warning Onset Time (WOT) measures how much reaction time the driver would have to prevent the outside tire from crossing a virtual 3-foot shoulder. The larger the WOT the better, as a larger WOT means the driver has more time to react to a warning and prevent an accident. The data I collected did not have shoulder width information. Therefore, I use a value of 0.91m (3.0ft). While this is narrower than the average shoulder width of U.S. highways, using this shoulder width combined with acceptable nuisance alarm rates lead to warning onset times which are just above the threshold of human reaction time, making this an interesting area to explore. When dealing with lane changes, this shoulder is actually a virtual shoulder, and is used as a baseline for computing the WOT. An alternative method would be to compute the WOT relative to when the vehicle actually crosses the virtual lane boundary. However, WOT would then be dependent on the selection of the virtual lane boundary, which would make comparing the results of different FOD parameters difficult.

In principle, I calculate the WOT as follows: When an alarm triggers, search forward in the data up to a certain point (usually about 3 seconds or so) and see if a lane change is found. If a lane change is found, then find the point where the vehicle's outside wheel first exceeds the (virtual) shoulder width. There are now two points: the alarm trigger point, and the shoulder excursion point. The difference in time between these two points is the WOT.

In practice, it is not possible to find the exact shoulder excursion point, due to the way Autotrak handles lane changes. Figure 3-12 illustrates this. The top plot shows the lateral position during the six seconds surrounding a right lane change, and the bottom plot is the lateral velocity during this same period. Positive lateral position and velocity are rightward, and negative are leftward. The lateral position indicates the position of the center of the vehicle relative to the center of the lane. Notice the discontinuity near T=1327. This marks a lane change event, and is the point where Autotrak locks onto the new lane. In general, the data, particularly the lateral velocity, for two to three seconds after this point is unreliable, and cannot be used.

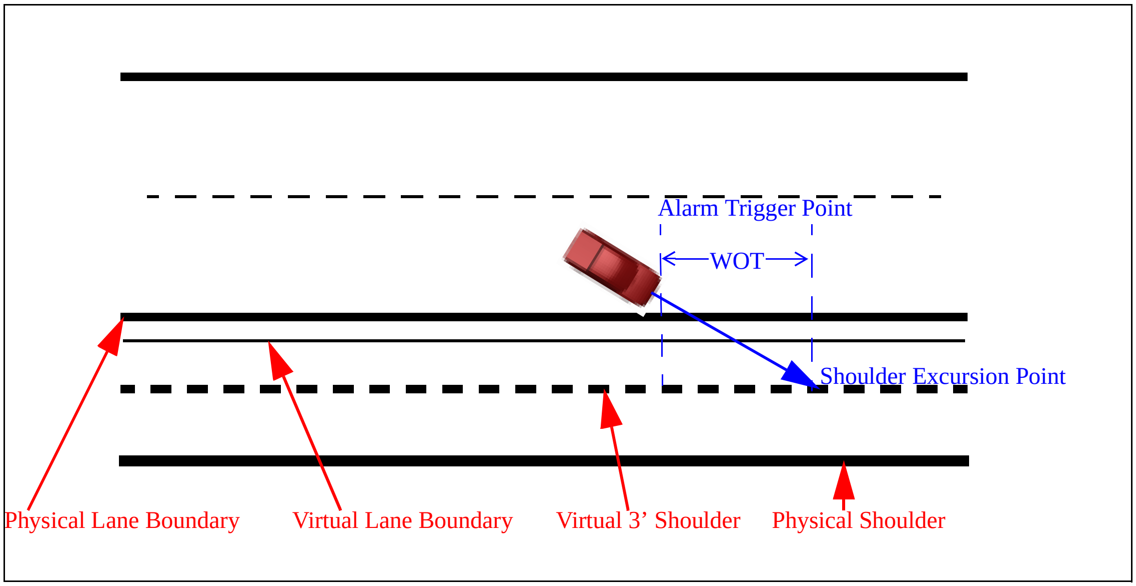

Figure 3-11: Depiction of Warning Onset Time Calculation.

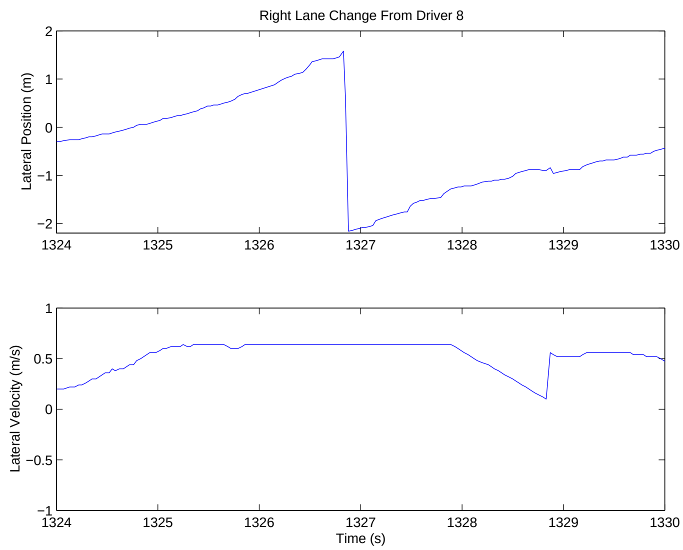

Figure 3-12: Right Lane Change. The discontinuitiy is a lane change event, and marks where Autotrak locks onto the new lane.

Given a 0.91m virtual shoulder, a 3.6m lane width, and a 1.8m vehicle width, the shoulder excursion point occurs when the center of the vehicle is 1.81m away from the lane center. However, the lane change event occurs before the vehicle is actually at the shoulder excursion point. At the lane change event, the lateral position is around 1.55m. To accommodate this, I linearly extrapolate lateral position to find the time of the shoulder excursion point.

The bottom plot on Figure 3-12 shows the actual lateral velocity (which, again, is valid only until the lane change event). The lateral velocity is constant for nearly two seconds before the lane change event, indicating that this sort of linear extrapolation of lateral position is justified. This is another reason for not using larger shoulder widths. The extrapolation required to find the shoulder excursion point for a six or nine foot shoulder would be significant, and could affect the WOT calculation.

Ideally, the WOT should be equal to the FOD lookahead time, i.e., if we are trying to predict lane changes 1.0 seconds in advance, then when we *do* predict a lane change, the prediction should occur 1.0 seconds before the event. However, the FOD algorithm predicts when the vehicle will be beyond the virtual lane boundary. The WOT is calculated relative to the 0.91m virtual shoulder. This discrepancy means that the WOT will not be equal to the prediction timestep.

There is a secondary effect which causes variance in the WOT calculation. The linearity in Figure 3-12 is not completely consistent. It is possible to find lane changes where the linear portion is of shorter duration, and there are cases where the lateral velocity is not constant between the alarm trigger point and the shoulder excursion point. This variability affects the warning onset time.

3.4.3.4 Nuisance Alarm Rate Calculation

The Nuisance Alarm Rate (NAR) is the other metric which I use to calculate alarm system performance. The Nuisance Alarm Rate indicates the number of alarms per hour which are not true alarms. This is an important metric because if the NAR is too high, then drivers might become frustrated with the system, and turn it off, or become habituated, and not react in time to a true alarm. A driver's NAR is calculated by counting all the alarms which do not correspond to lane changes, and dividing by the length of the data (in hours). I do not attempt to further classify which nuisance alarm are false alarms, based on incorrect sensor information or model inaccuracy, as defined in Section 1.2. The naturalistic data has a confidence measure which allows data which the AutoTrak system believes to be inaccurate to be filtered out.

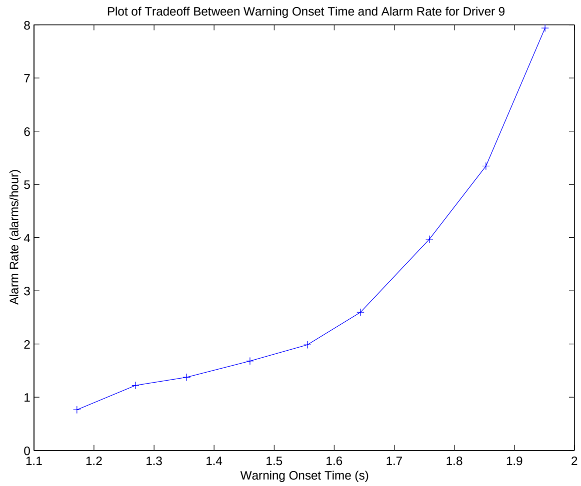

Figure 3-13: Illustration of trade-off between warning onset time and nuisance alarm

3.4.3.5 Warning Onset Time vs. Nuisance Alarm Rate Trade-off

There is a natural trade-off between warning onset time and nuisance alarm rate. Namely, as the warning onset time increases, so does the nuisance alarm rate. This is because, to increase warning onset time, either the lookahead timestep, *T*, has to be increased, or the virtual lane boundary, *V*, has to be decreased. Either change makes the algorithm more sensitive, and causes more triggers in areas which are not lane changes. Figure 3-13 shows the trade-off for D(9). The x-axis is warning onset time, and the y-axis is minimum nuisance alarm rate achievable for the corresponding warning onset time. The trade-off is not linear -- nuisance alarm rate tends to increase fairly quickly after a certain point. This is the main constraint which will affect all the experiments presented below, and the performance of the FOD algorithm.

3.5 Methodology

The above sections described the evaluation metrics which were used. Here, I give details of the methods which are evaluated and how the FOD training and testing is done.

3.5.1 The Warning Algorithms

This section describes the different methods which are evaluated: roadside rumble strips, TLC, AutoTrak FOD, Generic Learned (GL), and Individual Learned (IL). The infrastructure-based roadside rumble strips system is simulated to have a warning threshold set 0.15m beyond the physical lane boundary. The TLC model uses a 1.0s lookahead time, and triggers when the vehicle is predicted to cross the physical lane boundary. The AutoTrak FOD model uses a lookahead time of 0.85s and a virtual lane boundary of 0.10m. These parameters were hand selected to trigger just before the tire crosses the lane boundary during lane changes and fast lane departures, and to allow the driver to approach the lane boundary at low lateral velocities without a trigger.

The Generic Learned model attempts to answer the question: “What would the performance be if we found one set of FOD parameters (lookahead time and virtual lane boundary) for a set of drivers, and tested these parameters on a new driver?”. Therefore, when training a generic model for D(n), the data for all drivers except that of D(n) is used. Indirectly, the results of this model provide information on how similar driver performance is. The Individual Learned model, as the name implies, attempts to find the best FOD parameters for a particular driver, ignoring all other drivers.

3.5.2 Roadside Rumble Strips Results

TABLE 3-3 shows how the roadside rumble strips system would perform on the naturalistic datasets. In general, the nuisance alarm rate is low, except for D(1). D(1)’s high NAR (greater than 5/hour) implies that D(1) frequently exceeds the physical lane boundary by 0.15m during normal driving. The NAR for the other drivers is low, and most likely acceptable. However, the warning onset time is also low. In all cases, it is less than 1 second. While this is within the boundaries of human reaction time, as shown by [REF], it is not much time in which to react to a dangerous situation. As [75] shows, the roadside rumble strips system has reduced run off road accidents by 70%. A system which provided more warning onset time would presumably be even more effective on a roadway with wide shoulders and even work on roadways with narrow shoulders.

Driver	WOT (s)	NAR (alarms/ hour)	Nuisance Alarms
D(1)	0.99	5.36	28
D(5)	0.88	0.36	1
D(7)	0.96	1.39	2
D(8)	0.91	0.00	0
D(9)	0.94	1.07	7

TABLE 3-3 : Roadside Rumble Strips Warning Results

Driver	WOT (s)	NAR (Alarms/ Hour)	Nuisance Alarms
D(1)	1.87	41.90	219
D(5)	1.59	6.88	19
D(7)	1.78	22.22	32
D(8)	1.74	10.99	28
D(9)	1.79	5.35	35

TABLE 3-4 : TLC Warning Results

3.5.3 TLC Results

While the roadside rumble strips system has an nuisance alarm rate which is acceptable for most drivers, the warning onset time is low. TLC can provide greater warning onset time. However, as TABLE 3-4 shows, this comes at the expense of additional nuisance alarms. The warning onset time increases by over 80%. The NAR rises dramatically, though. For D(1), it is an alarm every 1.5 minutes, which is clearly unacceptable. The lowest alarm rate (5.35/hour for D(9)) is equal to the highest alarm rate generated by the roadside rumble strips system.

3.5.4 FOD Training

In the previous two sections, I present the results of two current warning systems, roadside rumble strips and TLC, as tested on real world driving data. The results show that both systems have advantages; roadside rumble strips has few nuisance alarms (but a low warning onset time), while TLC has a high warning onset time (but high nuisance alarm rate). In this section, I present the results of experiments using an FOD warning algorithm, and show that this approach provides more warning onset time than roadside rumble strips, while maintaining a low nuisance alarm rate. In this first section, I give the details of how the FOD warning algorithm can be trained for a particular driver (the individually learned model), or for all drivers (the generic learned model). These two cases are described in Section 3.5.1.

The two metrics which are used to compare warning system performance, warning onset time and nuisance alarm rate, do not have the same units, and cannot be directly compared. That is, if one alarm system produces a 1.5s WOT and a 0.5 NAR, is that better than the system which gives 1.75s WOT but a higher 0.75 NAR? It is possible to define acceptable trade-offs using decision theory, as done by Goodrich [25]. However, any such definition would be biased by the interpretation of the experimenter. Instead, I try to compare warning systems in which one of the metrics is equivalent. It is much easier to believe that a warning system with a 1.5s WOT and a 0.5 NAR is better than one with a 1.5s WOT but a 2.0 NAR.

To compare the generic models against AutoTrak, *FOD(0.85, 0.1, n*) is computed, $n \in (1, 5, 7, 8, 9)$. Taking Driver 1 as an example, *FOD(0.85, 0.1, 1)* results in a specific WOT, *w*, and NAR, *a*. Then, *FOD(Dgen(1), w)* is found, which uses all the drivers *except D(1)* and finds a set of parameters which gives the lowest nuisance alarm rate at a WOT of *w*. These parameters are then applied to D(1), and the resulting nuisance alarm rate is compared against the nuisance alarm rate of *FOD(0.85, 0.1)*. The parameter search is done using brute force, as the parameter space (*T* and *V*) is relatively small. The algorithm, using D(1) as an example, is:

- Apply each FOD parameter pair to all the data in *Dgen(1)*, and store the WOT and NAR results
- Find the set of FOD parameters, *p*, which results in the desired WOT.
- Find the parameters in *p* with the lowest NAR.

- Apply this parameter to D(1), and store the WOT and NAR results.

This can be considered to be an optimization problem to satisfy the following constraints:

$$WOT = wot(FOD(0.85, 0.10)) \tag{3-4}$$

$$NAR = MIN\left(\left(\sum_{T = 0.0}^{8.0} \sum_{V = 0.0}^{0.9} nar(FOD(T, V))\right)\right) \tag{3-5}$$

The individual model is trained in the same way, except that a leave-one-out cross validation is used for subsets of the data for a particular driver. The desired WOT is computed using *FOD(0.85, 0.1, 1)* (for D(1)). The data for D(1) is split into approximately half-hour segments. Training is done as follows:

- Pick one of the half-hour segments to use as a test set
- Do a brute-force FOD parameter search and find the parameters which produce the desired WOT and minimize the nuisance alarm rate on the remaining data (the training set), using the algorithm above.
- Apply this parameter to the test set, and store the WOT and NAR.
- Repeat, using a different segment as the new test set, until all segments have been tested.

This results in *n* estimates of WOT, and NAR, where *n* is the number of half-hour segments for the driver. These *n* estimates are then averaged together.

3.5.5 FOD Results

TABLE 3-5 shows the results of comparisons against the AutoTrak FOD parameters. In general, the learned models performed better than the AutoTrak parameters. In most cases, the generic model performed better than the individually learned model, although usually by two or fewer alarms. With D(7), however, the difference in number of alarms between the generic and individual models was 4. However, this difference comes at the expense of 0.12s of WOT. Furthermore, there are only about 1.4 hours of data on D(7). Of that, in the cross-validation, a half hour was used as test data, leaving only one hour of training data. It is possible that the lack of training data caused the lower performance. In contrast, D(1) and D(9) had

Driver	ATrak WOT (s)	Generic WOT (s)	Ind. WOT (s)	Atrak NAR (alarms/ hour)	Generic NAR (alarms/ hour)	Ind. NAR (alarms/ hour)	ATrak Nuis. Alarms	Generic Nuis. Alarms	Ind. Nuis. Alarms
1	1.64	1.71	1.61	20.28	11.29	8.61	106	59	**45**
5	1.44	1.33	1.39	3.62	2.17	2.90	10	**6**	8
7	1.52	1.36	1.48	10.42	5.56	8.33	15	**8**	12
8	1.55	1.57	1.49	4.32	4.32	4.71	**11**	**11**	12
9	1.57	1.52	1.54	3.21	1.68	1.99	21	**11**	13

TABLE 3-5 : Warning System Performance Comparison Between AutoTrak, Generic Learned, and Individual Learned. Bold indicates lowest nuisance alarms.

about 4.5 and 5.5 hours of training data, respectively. Although the generic model performed best for D(9), it was only by two alarms. The generic parameters may perform better at times because they capture a wider variety of driver behavior, causing the chosen generic parameters to allow for more weaving, as is explained in Section 4.2.7.3. The WOT for the generic results are usually somewhat lower than the trained results, except for D(8).

3.6 Modelling Additional Driver Behavior

The alarm decision model described above, and the warning algorithm it is a part of, make two assumptions: 1) Driver behavior does not change with road geometry, and 2) Driver behavior does not change over time. However, both these assumptions do not hold, as I demonstrate. Accounting for behavior changes due to road geometry and time results in a more complete alarm decision model, which improves lane departure warning system performance. The first additional behavior I model is curve-cutting, or the tendency to drift toward the inside of a curve. The second is local adaptation, which allows the driver to change behavior over time.

3.6.1 Curve Cutting

Curve-cutting is a behavior in which drivers tend to shift toward the inside of a curve. There are physiological reasons for doing this -- [40] has shown that when people negotiate curves, their eyes tend to focus on the inside of the curve. Also, when traversing a curve, if the driver starts at the outside, shifts inwards, and then back out, the radius of curvature is maximized and lateral forces are minimized, resulting in a smoother drive. Figure 3-14 shows a

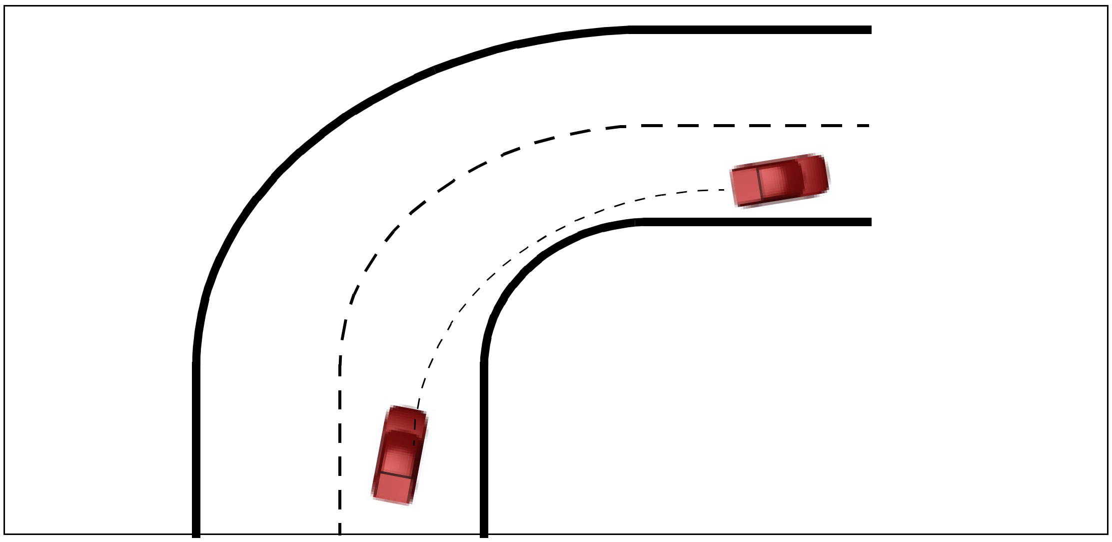

Figure 3-14: Illustration of Curve Cutting. Note how the vehicle moves towards the inside of the curve. This trajectory minimizes lateral forces.

depiction of curve cutting. Drivers do not always cut curves. Sometimes, they move to the outside of a curve, if they do not steer in time, or understeer.

3.6.1.1 Curve Cutting Evidence

The naturalistic data provides evidence for curve cutting, as shown in Figure 3-15. This figure shows the mean lateral position for five different cases: shallow and sharp left curving roads, straight roads, and shallow and sharp right curving roads. A shallow curve is defined as one where the radius of curvature is between 1000m and 2000m. A sharp curve is one where the radius of curvature is less than 1000m. A straight road is any whose radius of curvature is greater than 2000m. Negative lateral positions are left of center, and positive lateral positions are right of center.

The table shows that, in every case, the mean lateral position during left curves is to the left of the mean lateral position during straight segments. This effect is stronger for sharp curves than for shallow curves. The same is true for right curves, except for D(1), where there the straight mean is slightly rightward of the right curve mean. This could indicate that either D(1) cuts slightly to the outside of right curves, or does not cut at all. This indicates that in general, driver mean lane position varies inversely with radius of curvature -- the sharper the curve, the more curve cutting the driver exhibits.

3.6.1.2 Curve Cutting Model

Curve cutting behavior can cause additional false alarms, particularly if the driver exceeds the lane boundary. It is possible to account for this by temporarily increasing the virtual lane boundary on the direction toward the inside of the curve, allowing the driver more

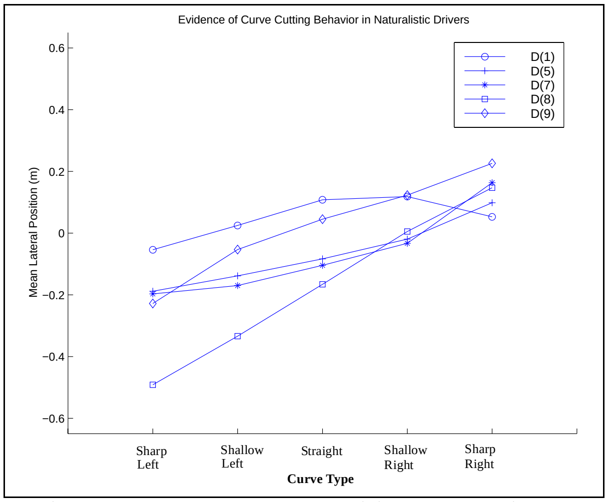

Figure 3-15: Mean Lane Position For Various Curve Types

leeway. The evidence in the previous section implied that curve cutting behavior varies inversely with radius of curvature. Therefore, the alarm decision model is updated to include a term for curve cutting, and an alarm is now triggered if the following relation is true:

$$L_p + TL_v > \left\{ V + \left(\frac{MIN\left(c\left(\left(\left(\frac{2000}{R}\right)\right), 0.50\right)\right)}{100} \right) \right\} \tag{3-6}$$

Where L_p is the lateral position (in meters), L_v is the lateral velocity, T is the FOD lookahead time, V is the FOD virtual lane boundary, c is a curve cutting weight, and R is the instantaneous radius of curvature (in meters). Note that this is an inversely proportional mapping from radius of curvature to virtual lane boundary extension. The curve cutting allowance is applied only if R is less than 2000m, and is capped at 0.50m. The capping is done so that extreme amounts of curve cutting are not allowed. The evidence in the previous section suggests that the change in mean lateral position between straight roads and sharp curves can be as high as 0.35m (for D(8)), so 0.50m is a reasonable value at which to cap the allowance. The 100 in the denominator converts the allowance from centimeters to meters. The weighing parameter, c, could be tailored to an individual driver; however, as the next section shows, the value is not critical.

3.6.1.3 Curve Cutting Parameter Analysis

To determine the sensitivity of warning system performance as the curve cutting factor, c, is varied, I did an experiment which tested different values of c over the range [0.0, 16.0], and compared the WOT and NAR. For each driver, FOD parameters were found which produced approximately five alarms/hour, while maximizing WOT, when *no* curve cutting allowance was used. Figure 3-16 shows the results. There are five graphs in the figure, one for each driver. The top plot in each is the WOT vs. curve cutting factor and the bottom plot shows NAR. In all cases, the WOT decreases as c increases. This makes sense, as the increasing c allows the driver to deviate further from the lane center, so lane changes towards the inside of a curve will be detected later.

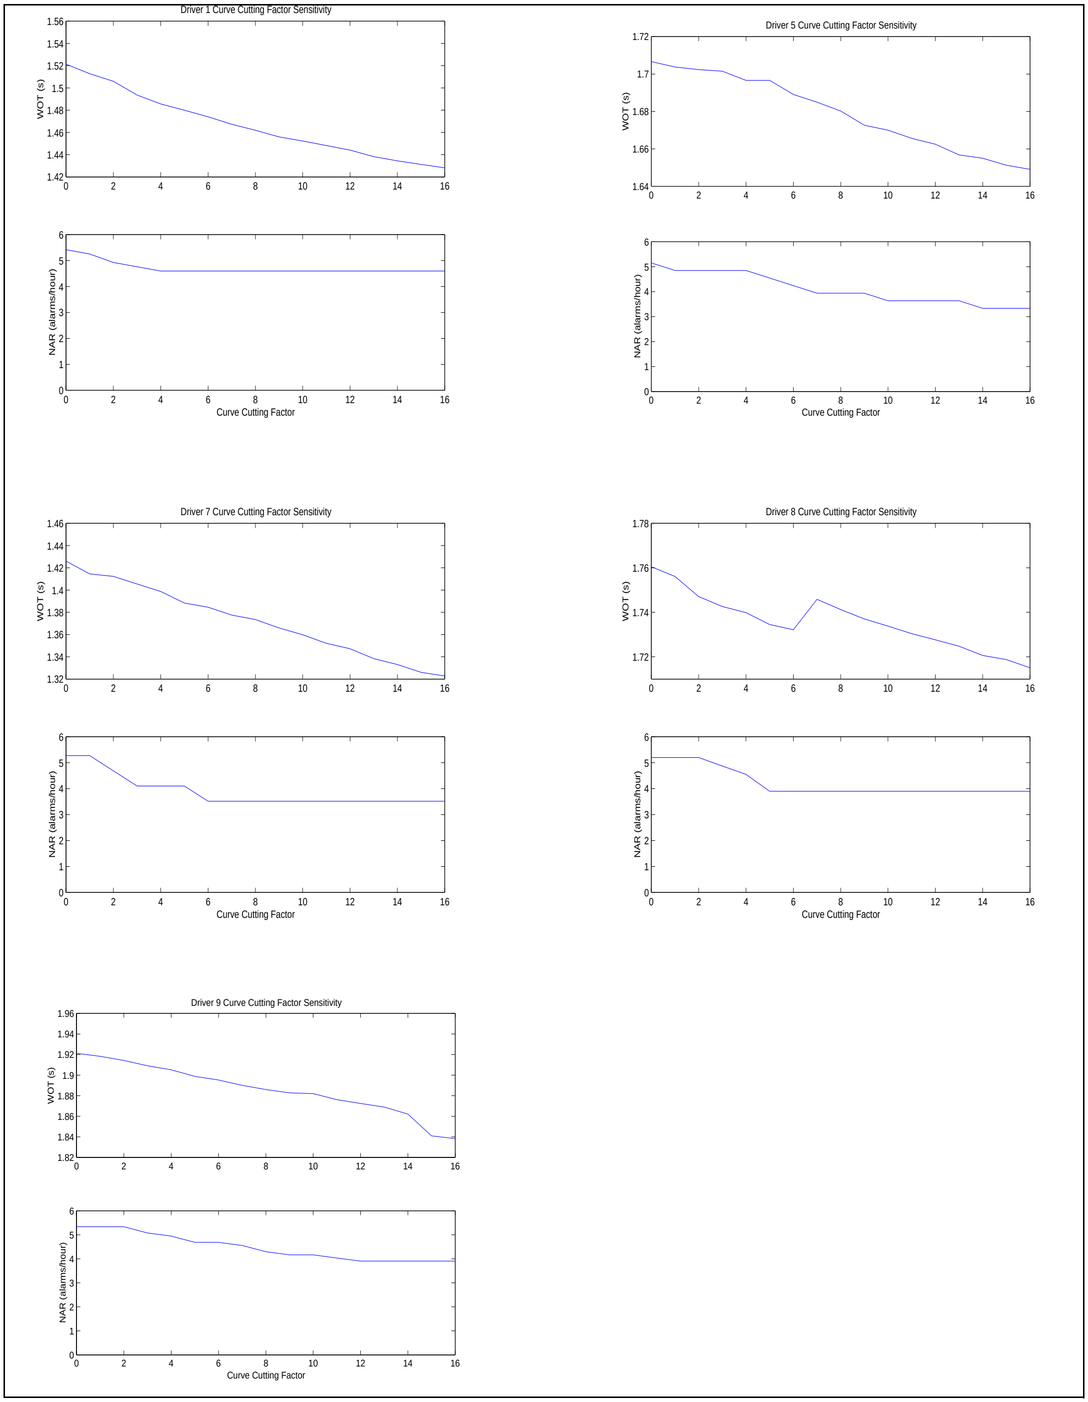

Figure 3-16: Results of Curve Cutting Factor Parameter Analysis. For each plot, the top window is WOT and the bottom window is NAR.

Driver	ATrak WOT (s)	Generic WOT (s)	Ind. WOT (s)	Atrak NAR (alarms/ hour)	Generic NAR (alarms/ hour)	Ind. NAR (alarms/ hour)	ATrak Nuis. Alarms	Generic Nuis. Alarms	Ind. Nuis Alarms
1	1.59	1.63	1.54	14.92	13.39	6.51	78	70	**34**
5	1.39	1.29	1.36	1.45	0.72	0.72	4	**2**	**2**
7	1.47	1.32	1.43	6.25	4.86	5.56	9	**7**	8
8	1.53	1.58	1.45	2.36	4.32	3.14	**6**	11	8
9	1.52	1.49	1.47	1.99	1.07	1.07	13	**7**	**7**

TABLE 3-6 : Warning System Performance Comparison Between AutoTrak, Generic Learned, and Individual Learned, using a Curve Cutting Allowance. Bold indicates lowest nuisance alarms.

The nuisance alarm rate drops with increasing c, as alarms due to curve cutting are eliminated. This decrease in nuisance alarm rate plateaus, indicating that not all the nuisance alarms are caused by curve cutting. Note that the decrease in nuisance alarm rate seems to plateau around the area c=8.0. Picking a value of c=8.0 seems reasonable, as at that point, the drop in WOT is less than 0.10s, and the nuisance alarm rate is mostly unchanging beyond that value. It is possible that allowing for curve cutting could have an adverse impact on WOT, if a lane departure were to occur towards the inside of a sharp curve. In that case, the curve cutting adaptation would result in the alarm triggering later, reducing the driver’s reaction time. However, based on an analysis of 200 ROR crashes, [60] shows that crashes where the vehicle drifts off the inside of a curve are four times less likely than accidents where the driver drifts off the outside of the curve. The next section shows the results of the same experiment described in Section 3.5, but with the additional curve cutting term in the model, and these results are discussed next.

3.6.1.4 Curve Cutting Modelling Results

The results against AutoTrak are shown in TABLE 3-6. The first thing to notice is that, across the board, the alarm rates have gone down, when compared against the results with no curve cutting allowance. AutoTrak does better with the addition of this term, along with the trained models. This reduction in nuisance alarm rate comes with almost no real change in

WOT. Again, the individual model for D(1) has a lower NAR than the generic and AutoTrak models. For the other drivers, either the individual or generic learned model does best, except for D(8), where the AutoTrak model has two fewer alarms than the trained model.

3.6.2 Local Adaptation

The training in the previous sections tries to find one set of FOD parameters which minimize nuisance alarm rate for a desired warning onset time. As we have seen, this approach is generally an improvement over the hand picked AutoTrak parameters. These trained parameters, while globally a good choice, may not be the best choice locally. That is, if a driver spends 98% of the time driving nearly down the center of the lane, without swerving much, then training may yield a narrower virtual lane boundary, as this would increase warning onset time. If, during the other two percent of the time, the driver is hugging one side of the lane, the narrow virtual lane boundary could generate false alarms. This shift in lateral position could occur for many reasons: construction zones, passing trucks, a narrow shoulder on the opposite side, driver behavior, etc.

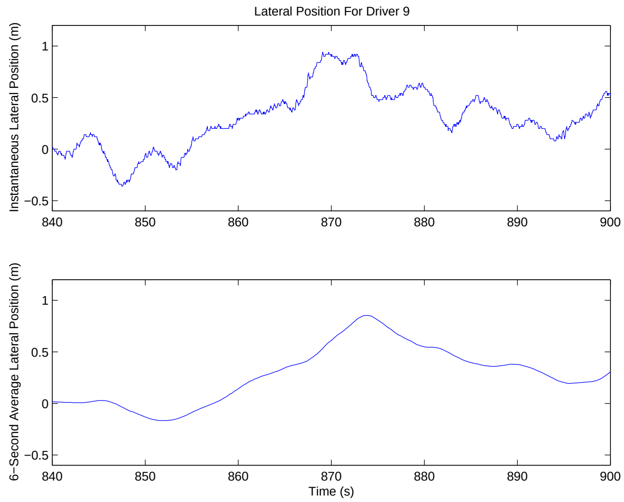

Figure 3-17: Evidence of local deviations in lane position.

3.6.2.1 Local Adaptation Evidence

Figure 3-17 shows an example of this, from one minute of D(9)'s data. The top plot shows one minute of instantaneous lateral position data, while the bottom plot is a six-second moving average of lateral position. Over the course of one minute, the driver has a lateral position shift of over one meter, with about 8 seconds spent nearly touching the right lane boundary (time ~868 - ~875s). Appendix C shows further evidence, using the Navlab 8 data, that drivers shift because of adjacent vehicles.

Ideally, it would be nice to know why the driver's behavior has changed. If a warning system could detect a passing truck, or knew the vehicle was entering a construction zone, then that information could be used to modify the warning algorithm. In practice, this type of information is not available. AutoTrak has neither obstacle detection capability, not any knowledge of the surrounding road conditions. While it is possible to add sensors around a vehicle to detect other objects, this is an expensive, non-portable solution. Similarly, there are currently no roadway information systems which transmit local road conditions, such as construction activity.

3.6.2.2 Local Adaptation Model

It is still possible to adapt to local deviations, however, by looking at a lagged (by the window size) moving average of the vehicle's lateral position. By doing so, a warning system would know *when* the driver's behavior has changed, although it could not infer the reason for that change. If we assume that all such short term changes are due to external circumstances or decisions made by the driver, rather than to inattention or drowsiness, the alarm decision model can be modified to allow the driver more latitude in these circumstances.

In doing this, there are two questions to consider: How fast should the warning algorithm adapt to the driver? and how should the virtual lane boundary be modified? One way is to trigger an alarm if the following is true:

$$L_p + TL_v > \{V + (a(\overline{L_{p(t-n,\,t)}}))\} \tag{3-7}$$

Where L_p is lateral position, L_v is lateral velocity, T is the lookahead time, and V is the virtual lane boundary. $\overline{L_{p(t,\,t-n)}}$ is an average of lateral position from time t to time t-n, where n is in seconds, and a is the local adaptation factor, which weights the contribution of the current mean lateral position to the virtual lane boundary. Therefore, there are two new parameters, n and a. As with the curve cutting factor described in Section 3.6.1.2, a parameter analysis of the new parameters shows that the selection of these parameters is not critical.

3.6.2.3 Local Adaptation Parameter Analysis.

This section describes a parameter analysis in which the parameters n and a from Equation (3-7) are varied over the range $a \in [0.0, 1.0]$ and $n \in [1.0, 20.0]$, using FOD and curve cutting parameters which produce approximately 10 alarms/hour. Figure 3-18 shows the warning onset time (on the left) and nuisance alarm rate (on the right) results. The x-axis is n, and the 11 lines on each graph represent different values of a, from 0.0 to 1.0, in increments of 0.1, starting at the top. The topmost line in each graph represents a=0.0, which is the default case where no local adaptation is added to the virtual lane boundary.

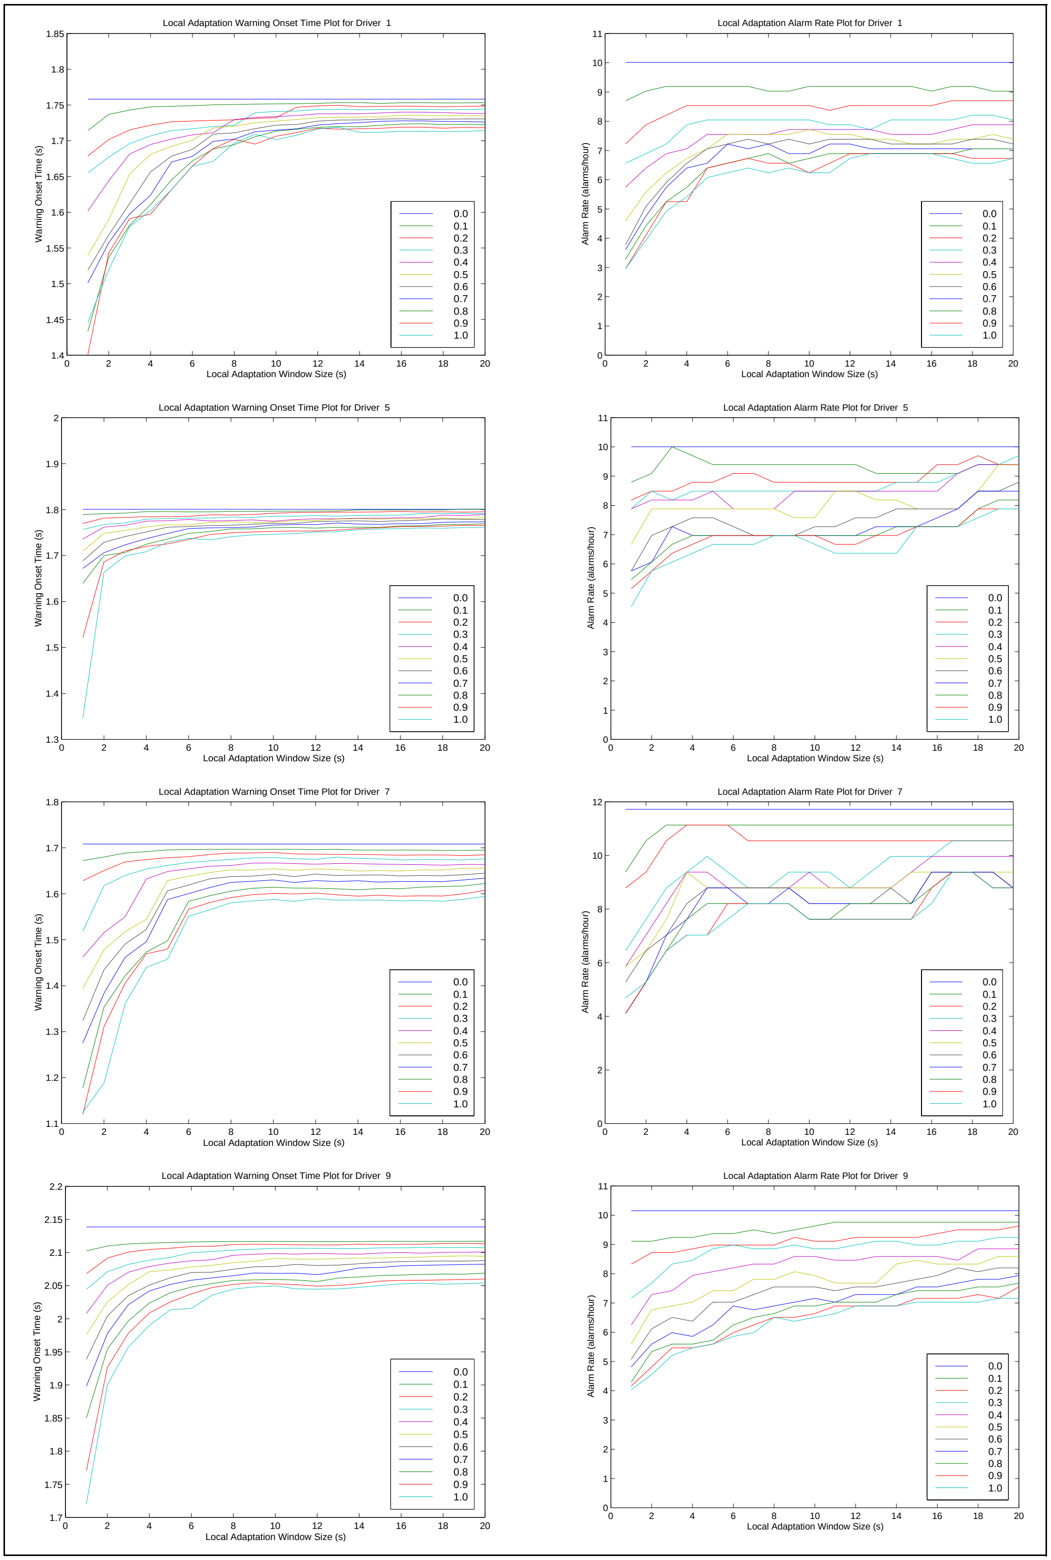

Figure 3-18: Local Adaptation Parameter Analysis. The left graphs are WOT results, and the right graphs are NAR results. Each graph contains 11 lines, where the top line is *a*=0.0, and the bottom line is *a*=1.0.

Driver	ATrak WOT (s)	Generic WOT (s)	Ind. WOT (s)	Atrak NAR (alarms/ hour)	Generic NAR (alarms/ hour)	Ind. NAR (alarms/ hour)	ATrak Alarms	Generic Alarms	Ind. Alarms
1	1.48	1.54	1.44	7.08	6.12	3.44	37	32	**18**
5	1.27	1.19	1.23	0.36	0.00	0.00	1	**0**	**0**
7	1.23	1.10	1.26	4.17	2.78	3.47	6	**4**	5
8	1.41	1.43	1.37	0.79	2.75	1.18	**2**	7	3
9	1.37	1.30	1.34	0.92	0.46	0.61	6	**3**	4

TABLE 3-7 : Warning System Performance Comparison Between AutoTrak, Generic Learned, and Individual Learned, in Alarms/Hour Format, using a Curve Cutting Allowance and Local Adaptation. Bold indicates the lowest number of nuisance alarms.

As *a* increases (for a fixed value of *n*), we notice two things: the warning onset time drops, and the nuisance alarm rate drops. This is because when *a* increases, the local average's contribution to the virtual lane boundary increases, pushing it further out, and decreasing the probability of a nuisance alarm. Similarly, since the virtual lane boundary is wider, it takes longer for true alarms to be detected, decreasing the warning onset time.

For low values of *n*, the virtual lane boundary adapts quickly to changes in lateral positon. At the extreme, when *n*=1.0, the nuisance alarm rate and warning onset time drop dramatically. This is because the system is adapting too quickly, so that alarms (both true and nuisance) become difficult to detect. At higher values of *n*, beyond 8.0s, there is not much of a change in warning onset time or nuisance alarm rate. In that region, the value of *a* has more of an effect. Based on these results, I decided to use *n*=6.0s and *a*=0.8.

3.6.2.4 Local Adaptation Results

Tables TABLE 3-7 shows the results of AutoTrak vs. generic vs. individual training performance, using an FOD model which includes both curve cutting and local adaptation. The addition of local adaptation lowered alarm rates for all the cases. As in the previous cases, using an individually trained model for D(1) greatly reduces the number of false alarms.

The results for the other drivers are harder to interpret. Either an individual or generic trained model has a lower nuisance alarm rate than AutoTrak, except for D(8). However, there is only a 1 alarm difference between AutoTrak's performance and the individually trained model. In all the cases, except for D(1), the difference in nuisance alarm rate between the individually trained model and AutoTrak is less than or equal to two. Overall, the number of alarms is becoming small enough that direct comparisons are difficult.

3.7 Results Analysis

The previous sections presented alarm system performance as the FOD alarm decision model grew to include more aspects of driver behavior. This sections ties together the results of the AutoTrak comparisons, over the three different models, and presents further data on the performance improvement due to modelling curve cutting and local adaptation behavior. Finally, I analyze the parameter choices made by the individual trained model, concentrating on D(1) and D(9), as these drivers had the most data.

3.7.1 Combined Results

Figure 3-19 shows a graphical summary of the experiments performed against AutoTrak. The top graph shows the warning onset time for all five subjects over different model combinations and training methods. The bottom graph shows the same for nuisance alarm rate. The first thing to notice is that the nuisance alarm rate for D(1) is generally higher than that of the other drivers. This is due to the D(1)'s higher lateral position standard deviation, from TABLE 3-2. D(1) is also the driver who is helped most by using learned parameters. However, there is a caveat to this result. As I discussed in Section 3.4.2, D(1)'s data was contaminated by another driver. While the subject claimed that the majority of the driving was done by him, there is the possibility that some of the D(1) data used in this analysis is actually that of another driver. If this is the case, it may be more accurate to claim that the differences seen in D(1)'s behavior are not solely due to individual characteristics. Rather, they may in some part be affected by weather conditions (D(1)'s trip was in winter, on occasionally snowy roads), or the type of vehicle being driven.

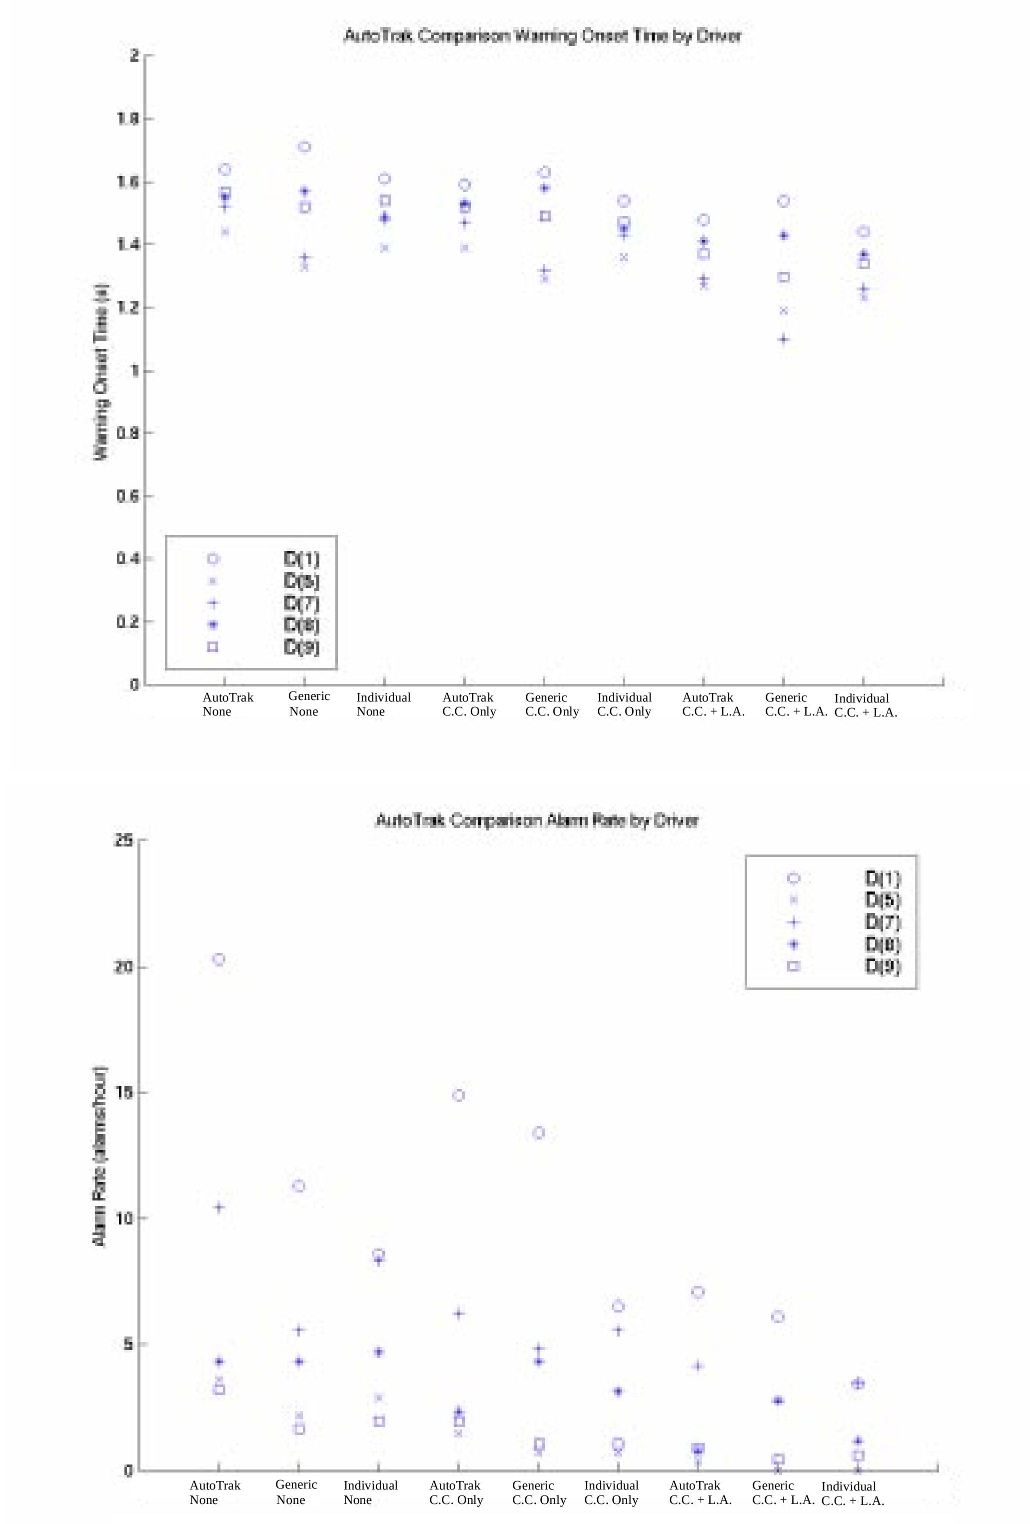

Figure 3-19: Combined results of the alarm performance comparison between the AutoTrak, Generic, and Individual models, over different combinations of alarm decision models.

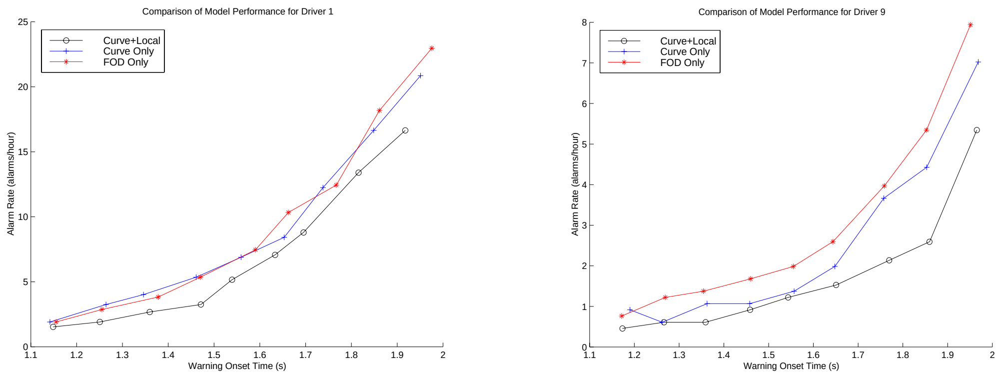

Figure 3-20: Comparison of Model Performance for Drivers 1 and 9.

Ideally, the warning onset times within each driver would be identical over the different combinations of training and models, allowing for a direct comparison of alarm rates. However, this is not the case. There is a correlation between improved nuisance alarm rate performance, and decreasing warning onset time. Therefore, it is possible that the extended models do not improve overall performance -- rather, they just slide performance into an area with lower WOT and NAR. Figure 3-20 shows the results of an experiment to determine whether this is happening. I computed the lowest possible nuisance alarm rate (using individual training) over a range of desired warning onset times. That is, “For a warning onset time of 1.2s, what is the lowest achievable nuisance alarm rate using different alarm decision models?”. For any particular WOT, using a combination of curve cutting and local adaptation should provide the lowest nuisance alarm rate. From Figure 3-20, we see this happens. The next lowest nuisance alarm rate should be from the curve cutting only case. While this is true for D(9), it doesn’t appear to hold for D(1), implying that modelling curve-cutting may not be very beneficial for this driver. There is further evidence for this in Figure 3-15, which shows D(1)’s curve cutting behavior. The mean lateral position for right curves is actually slightly left of the mean lateral position for straight roads, indicating that D(1) does not curve cut in that direction. Further, the curve cutting evidence figure, Figure 3-16, shows that D(1)’s nuisance alarm rate plateaus higher than the other drivers, which implies that including a curve cutting term in the alarm decision model was not able to eliminate as many nuisance alarms for D(1) as for the others.

D(8) is the only driver for whom the AutoTrak parameters ever do better than either the generic or individual trained parameter. When using the curve cutting and local adaptation heuristics, AutoTrak has one fewer alarm than the individual model, over approximately 2.5 hours of data. When the differences are this small, it is difficult to claim that they are significant.

When compared against roadside rumble strips and TLC, the FOD results based on the full alarm decision model have a higher WOT than roadside rumble strips, while the alarm rate is similar. As [Wood94] shows, roadside rumble strips has reduced the run-off-road accident rate by about 70%. Presumably, a larger warning onset time would reduce the accident rate even further. In particular, D(1) is helped significantly by the trained FOD model over roadside rumble strips. While the FOD WOTs are not as high as those provided by TLC, the NARs are significantly lower. This shows that the new alarm decision model combines the advantages of both roadside rumble strips and TLC to improve overall warning system efficiency.

3.7.2 Parameter Analysis

This section looks at the parameters chosen by the individual training algorithm, and compares them against AutoTrak. In particular, I focus on D(1), where the difference between AutoTrak and trained performance is significant, and on D(9), where the difference is not so large. From TABLE 3-7, we can see that the AutoTrak parameters on D(1) resulted in WOT = 1.48s and NAR = 7.08, whereas the trained parameters led to WOT = 1.44s, and NAR = 3.44 alarms/hour. This is a 50% reduction in nuisance alarm rate, with no real loss in warning onset time.

Equation (3-3) showed that there are an infinite number of FOD parameters which provide the same warning onset time, assuming a constant lateral velocity during deviations. The left-most plots in Figure 3-21 show the AutoTrak parameters, along with the set of parameters which came out of the leave-one-out individual training method described in Section 3.5.4. The training algorithm selects these parameters to match AutoTrak's WOT, while minimizing NAR. The first thing to notice is that all the parameters lie on a straight line. This corresponds with Equation (3-3), which suggests that parameters which provide equal WOT

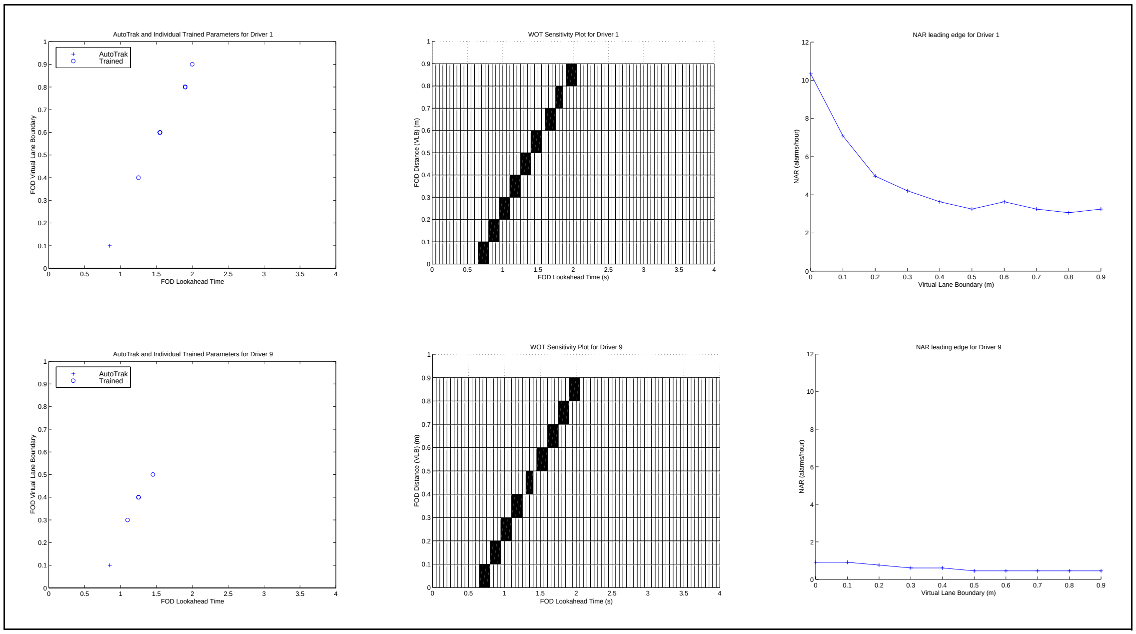

Figure 3-21: Parameter plots. The left plot shows the AutoTrak parameters along with the learned parameters. The second plot shows ALL parameters which provide a WOT equal to AutoTrak's performance. The third plot shows the nuisance alarm rate along the parameters of the second plot. The top row is D(1), the bottom row is D(9)

performance lie on a line. The middle plot, which is a plot of *all* the parameters which result in a WOT of 1.44s +/- 0.05s confirms this. The right-most plot shows the nuisance alarm rate along the line of equal WOTs. The x-axis of this plot corresponds to different values of the virtual lane boundary, *V*, along the line of equal WOTs. From this, we can see that when *V* is low, the nuisance alarm rate is high, and that it decreases as *V* increases.

The second observation is that the learned parameters lie above and to the right of the AutoTrak parameters in an area where the lookahead time and virtual lane boundary are both larger. The reason for this is D(1)'s greater lateral position variance. D(1) tends to cross the lane boundary during normal driving, which means that the AutoTrak parameters, which have a virtual lane boundary setting of 0.1m, generate nuisance alarms. D(1) exceeds the lane boundary by 0.1m an average of 5 times per hour. This is reflected in the 7.08 alarms/hour nuisance alarm rate which the AutoTrak parameters produce.

A similar trend is observable in the bottom row of plots, which correspond to D(9). Here, the learned parameters have higher *V* values than AutoTrak, but not as high as D(1). The nuisance alarm rate along the line of equal WOTs for D(9) is very low, regardless of the value of *V*. This implies that D(9) is a tighter driver than D(1), and this is confirmed by the lateral position distributions and variances which were presented in Figure 3-6.

The correspondence between Equation (3-3),which shows that linear changes in FOD parameters combinations result in equivalent WOT, and the results in Figure 3-21, which show this linear relationship using real data, indicate that the underlying assumption regarding Equation (3-3) is valid. Namely, that during the portion of lane changes which the FOD model depends upon (namely, between the trigger point and shoulder excursion point), lateral velocity is relatively constant. This implies that, when drivers change lanes, they tend to make an initial steering input to start the vehicle on its trajectory, allow it to cross the lane, and then counter-steer to stabilize the vehicle in the new lane. This is different from Godthelp's [24] model, in which the driver executes a continuous sine-wave steering command. A random manual examination of lateral position and lateral velocity reveals that vehicles do *not* have a constant lateral velocity over the *entire* lane change maneuver -- rather, the lateral velocity increases until 1-2 seconds before the AutoTrak lane change event occurs, which is normally when the edge of the vehicle is about 0.5-0.8m beyond the lane boundary. This implies that, in general, true lane changes are really somewhere in between Godthelp's sinusoidal model and a pure impulsive model. However, Figure 3-21 suggests that at the alarm trigger point the lateral velocity is relatively constant. The average lateral velocity at the trigger point is, in fact, the slope of the line in the middle plot of Figure 3-21. For D(1), the slope of the line is 0.72m/s. It is the same for D(9). Perhaps this is not so surprising, as most people presumably have a range of comfortable lateral accelerations that they are willing to tolerate. Given the relatively short distances which are laterally traversed during a lane change, there is not much time to accelerate to high lateral velocity.

3.8 Summary

In this chapter, I present a Future Offset Distance (FOD) based alarm decision model, along with a performance analysis on real world data taken from five drivers, collected during the naturalistic data study, and compare this performance against two existing warning systems: roadside rumble strips and TLC. The FOD algorithm uses a linear prediction of future lane position to determine when the driver is about to depart the lane. I present the results of experiments on learning FOD parameters, and then add additional aspects of driver behavior such as curve cutting and local, short term deviation modelling to the alarm decision model, and show that this model, when applied to a lane departure warning system, reduces the number of nuisance alarms. I did a parameter analysis of these two additional factors, which showed that there is a range of parameters which perform similarly, and also showed the effects of extreme parameter values. Finally, a parameter analysis demonstrated that the lateral velocity near the alarm trigger points is relatively constant.

Overall, this chapter has shown that learning FOD parameters can help in certain cases, particularly when the driver has a large lateral position variance, and crosses the lane boundary often, as D(1). For the rest of the drivers, the benefits of learning FOD parameters is not conclusively demonstrated. These drivers all had relatively low lateral position variance, and did not exceed the lane boundary as often as D(1). While there does seem to be an advantage of either the generic learned model or the individual learned model over the AutoTrak for these other drivers, this advantage is usually quite small, and is not consistent, which makes significance difficult to judge. One issue which I believe affected performance was the amount of data available. While I collected over 70 hours of data during the naturalistic data study, only about 20 hours of data ended up being usable. Having 7 or 8 hours of data for each driver might have produced a more conclusive answer regarding variability in driver behavior. In Section 5.2, I present some lessons learned from this study, and suggestions for follow on studies.

CHAPTER 4 Driver Analysis

4.1 Introduction

In this chapter, I introduce the use of Memory Based Learning (MBL) for use as a driver behavior visualization and comparison tool. While most analyses of driver behavior involve looking at gross statistics such as means and standard deviations of lane keeping performance and steering wheel control, MBL allows one to examine driver behavior in specified situations. In particular, it is now possible to ask questions such as “Where is the driver likely to be in 1 second if he/she is currently 10cm left of center and has a lateral velocity of 10cm/s to the left?” This is important because it allows one to identify situations in which drivers differ from each other, which could then be exploited by a warning system, or provide insight into warning system performance and driver behavior.

The warning system results in the previous chapter give rise to two questions: 1) Why does learning FOD parameters produce better results for D(1), and 2) Why is D(1)’s overall performance so much worse than the other drivers? The answers to both these questions involve looking at how the FOD warning algorithm behaves in different parts of the vehicle state space (lateral position and velocity), and how the FOD parameter selection affects warning system performance. MBL, combined with probability measures and information theory, provides tools which can be used to answer the above questions.

The FOD warning algorithm I present in the previous chapter assumes that drivers have a constant lateral velocity over the lookahead period. If this were true, then the actual future lateral position of a vehicle would take a single value and be completely predictable. However, this is not true. There is uncertainty associated with a kinematic prediction, due to violations of the assumption of constant lateral velocity.

This uncertainty manifests itself through two main effects. The first and most common effect is that, for a given vehicle state, the actual future lateral position does not take on a single value. Rather, it takes on a distribution of values. This distribution is usually normal, with the standard deviation related to FOD lookahead time. The further ahead one tries to predict future vehicle state, the larger the standard deviation, indicating increasing uncertainty about the prediction.

The second effect is more interesting. I present results which show that the distribution of *actual* (not predicted) future lateral positions of a vehicle in a particular state can be bimodal, and that this bi-modality has an impact on warning system performance. For example, when making lane changes, there is a period before the lane change during which drivers could do one of two things: 1) Continue to drift away from the lane if their intent is to actually make a lane change, or 2) Correct their steering, and come back to the center of the lane. Predicting a lane change in these situations can increase the nuisance alarm rate. This bimodality in future vehicle state prediction presents a fundamental barrier to the performance of a warning algorithm based on just vehicle lateral position and lateral velocity. If a warning system is too sensitive, then it triggers warnings when drivers may actually end up correcting their behavior, which increases the number of nuisance alarms. This situation could be helped with the use of lateral acceleration information, which would allow earlier detection of the driver’s intent to continue drifting of begin correcting. In practice, this moves the problem of bimodal states earlier, but would not eliminate the effect.

This bimodality, and how it affects warning system performance, can be quantified using measures of uncertainty taken from information theory, along with false alarm probability calculations. False alarms, remember, are alarms caused by an error in vehicle state prediction or sensor area, and are a subset of nuisance alarms, as defined in Section 1.2. I calculate false alarm probabilities rather than nuisance alarm probabilities for two reasons: 1) To get at the heart of the appropriateness of using a 1st order FOD model, and 2) To avoid the issue of whether a given nuisance alarm would really be considered a nuisance in an online system, based on the definition in Section 1.2. As false alarms are a subset of nuisance alarms, any increase in false alarm rate results in an increase in nuisance alarm rate. Therefore, although I often discuss results in terms of their effects on nuisance alarm rate, calculations of false alarm probability are used to justify the results.

I use the uncertainty and false alarm probability calculations to show that the likelihood of triggering an alarm using trained parameters for D(1) is much less than the probability using the AutoTrak parameters. Since the overall probability of triggering an alarm is lower, the number of nuisance alarms is also lower. Similarly, I show that the behavior of D(1), who has a high nuisance alarm rate, contains less information, and is therefore fundamentally less predictable than the behavior D(9), who has a low nuisance alarm rate. This leads to better warning system performance for D(9), and poorer performance for D(1).

4.2 Memory Based Learning

In this section, I describe the use of Memory Based Learning, starting with an introduction to the basics of MBL, along with details on how I have used it in this work, followed by an example of system output.

4.2.1 Introduction

Memory Based Learning (MBL) [4] is a machine learning technique which is similar to techniques such as k-nearest-neighbor and locally weighted regression. MBL is most often used as a predictor or classifier, where the answer to a query is a function of answers to past training examples which are similar to the current query. This formulation requires that all past training examples be stored in memory, with a trade-off in decreased query time. The method used to store training examples often depends on the number of states in the problem. When the number of states is low (<= three), arrays or basic tree structures may be used. For higher dimensional state spaces, more advanced techniques such as KD-Trees [48] are used, to decrease query time.

For this work, the number of states is two: vehicle lateral position, and lateral velocity. The input (or query) is a vehicle state vector. The output is a distribution of actual future lateral positions, based on the past training data. The main idea is that instead of using a model (such as a 1st or 2nd order kinematic model) to *predict* the future position of the vehicle, use a lookup table to store the *actual* future lane position of the vehicle. Examining this output can provide insight into driver behavior.

4.2.2 Training Method

Because of the low dimensionality of the state input, I can represent the memory table as a 2-D array, where one axis is lateral position, and the other axis is lateral velocity. The axes are discretized into bins. I chose to use 0.05m for lateral position and 0.05m/s for lateral velocity. Using a smaller discretization does not make sense, due to the levels of noise in the naturalistic data. Using a larger discretization throws away information. The memory table is created as follows: For each training point p, which comes from the naturalistic training data, add the *actual* future lateral position for a particular lookahead time (by jumping forward in the data), L_p' to the appropriate bin. Iterate over all the training data. Figure 4-1 illustrates this. After training, each bin, b, contains a list of lateral positions, which are the distribution of future lateral position for the input state represented by b.

4.2.3 Memory Table Contents

Figure 4-2 shows the memory table after training on 2 hours of data from D(9). The x-axis is lateral position, and the y-axis is lateral velocity. Darker areas indicate higher density of training data. As expected, the central area is darkest, indicating that the driver spends most time within a half meter of the lane center, and with lateral velocities less than 0.5m/s. There are two areas circled in the figure -- these are areas of the state space which correspond to lane changes. For example, the trajectories of left lane changes extend from the central area towards the lower left portion of the state space, where both lateral position and lateral velocity are leftward. This correspond to a series of positions increasingly further to the left. The lane changes mostly have little variance in the lateral velocity axis, indicating that after a cer-

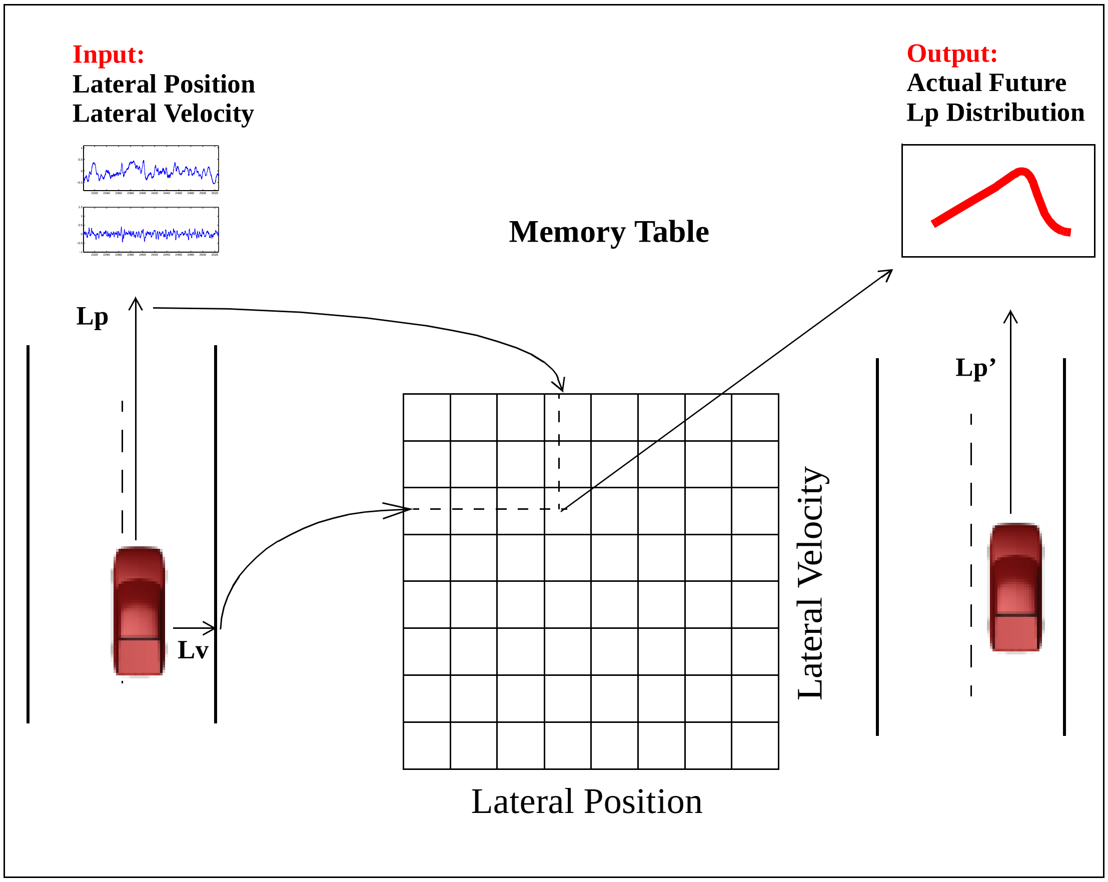

Figure 4-1: Memory Based Learning Training and Output. The input is vehicle state, the output is a distribution of actual future lane positions.

tain point, lane changes are roughly linear. Most of the lane changes occur at lateral velocities between 0.5 and 1.0m/s, which is in line with the 0.7m/s calculation for D(1) and D(9) given in Section 3.7.2.

4.2.4 Memory Table Density

As more training data is added, the state table becomes larger -- both in terms of the amount of data it contains, but also in the amount of the state space which is covered. This is because, as more and more data for a particular driver is added, the diversity of the state space which the vehicle has been through increases. Figure 4-3 shows this. The top left plot shows state space coverage after about one hour of training data, while the bottom center plot shows coverage after about five hours of data. The non-central areas of the state space have now been more fully explored. There is very little difference between the four and five hour case.

Figure 4-2: Memory table density after two hours of training. Areas in the state space which correspond to lane changes have been marked. The contrast for a cell shows the number of entries in the cell; the distribution of values within the cell is not shown.

This shows that 1-2 hours of data does not adequately represent all the behaviors that a driver may engage in. After about 4-5 hours of data, however, the density of the state space doesn't change as much, indicating a diminishing return.

4.2.5 Alarm Sensitivity

In this section, I show how the memory table can be used to visualize how changing the FOD parameters affects the triggering behavior of an FOD algorithm with 1st order vehicle state prediction, which I use in Section 3.3. Figure 4-4 shows a set of sensitivity plots. The x-axis is lateral position, and the y-axis is lateral velocity, which represents the state space of the 1st order FOD algorithm. The non-white areas are parts of the state space which have representation in the real data. The plots were generated using all of the nearly six hours of data

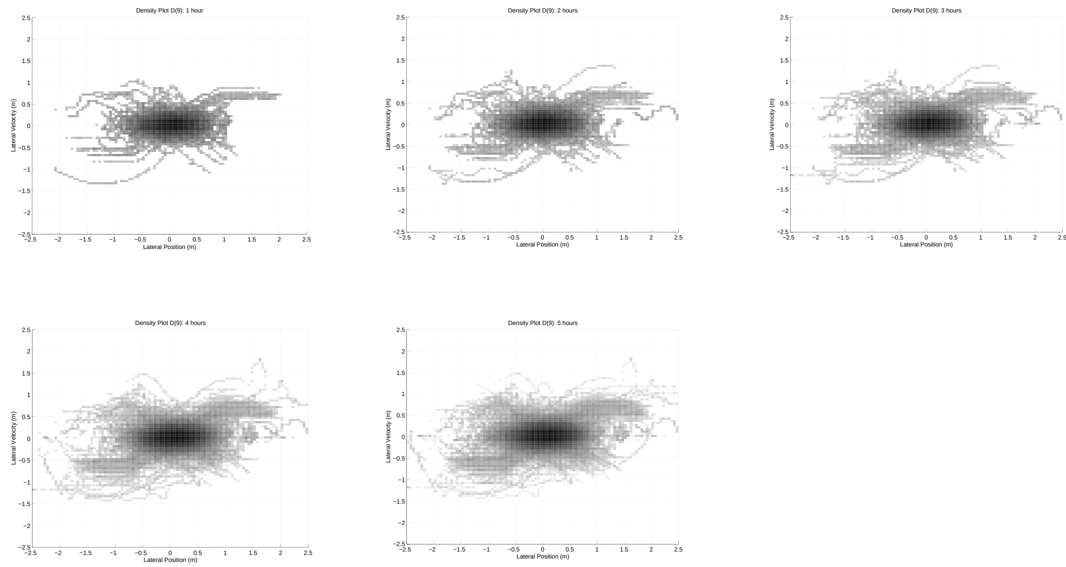

Figure 4-3: Memory table density as amount of training data increases. Darker areas indicate higher training density.

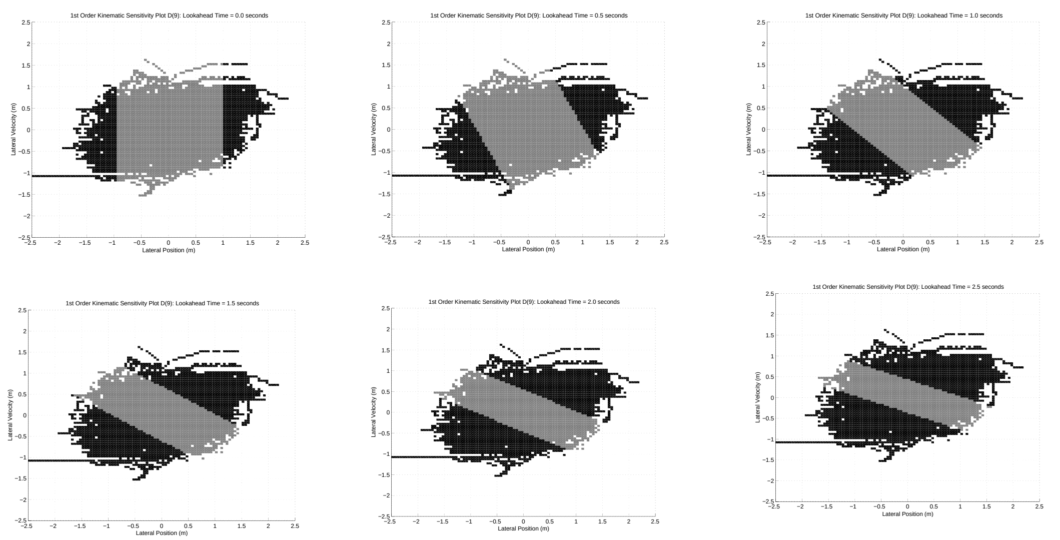

Figure 4-4: 1st Order Sensitivity Plots. FOD Virtual Lane Boundary is 10cm, while the FOD Lookahead Time ranges from 0.0s to 2.5s.

for D(1). The black areas are parts of the state space where a 1st order FOD algorithm would trigger and alarm, based on a 0.10m FOD virtual lane boundary, and an FOD lookahead time in the range of [0.0, 2.5]s. The gray areas represent states which do not trigger an alarm.

For all the plots, the central areas, in which the vehicle is near the center of the lane and has low lateral velocity, do not cause triggers. As we move towards more extreme lateral positions and velocities, the FOD algorithm begins to trigger alarms. The top-left plot shows the case when the lookahead time is 0.0s. Therefore, no actual prediction is done. An alarm triggers when the vehicle actually *is* 0.10m beyond the lane boundary. With an assumed 3.6m lane width and 1.8m vehicle width, this occurs when the vehicle's lateral position is 1.0m. Correspondingly, all the areas of the state space where the lateral position is greater than 1.0m show triggers.

As the FOD lookahead time changes, the trigger areas change shape. However, the boundary between trigger and non-trigger areas is always a straight line, as triggering is based on a linear prediction. The slope of the line is related to the lookahead time, and the y-intercept is related to the virtual lane boundary. Intuitively, this shows that as the lookahead time is increased (for a fixed virtual lane boundary), more and more parts of the state space cause triggers. If the lookahead time is increased too much, parts of the state space which are mostly related to normal driving cause triggers, which leads to an increase in the number of nuisance alarms. Similarly, if the lookahead time is too low, then triggers occur only in extreme circumstances, which reduces the warning onset time. However, as I show in the next section, even moderate lookahead times, such as used by AutoTrak, can cause triggers in areas of the state space where the right action is not predictable.

4.2.6 Sample Distributions

After training is done on a particular driver, it is possible to visualize the distribution of actual future lateral positions for a given input state. This will illustrate a fundamental limitation of any warning algorithm which is based solely on instantaneous lateral position and lateral velocity. Figure 4-5 shows three examples taken from D(9). Training was done to capture future lateral position with a lookahead time of 1.0 seconds. The title of each graph shows the current vehicle state while the corresponding histogram shows the distribution of actual future lateral position, one second in the future, for all the training data of that state. The dashed lines demarcate the physical lane boundary, assuming a vehicle width of 1.8m and a lane width of 3.6m. In the following examples, I assume that there is no virtual lane boundary.

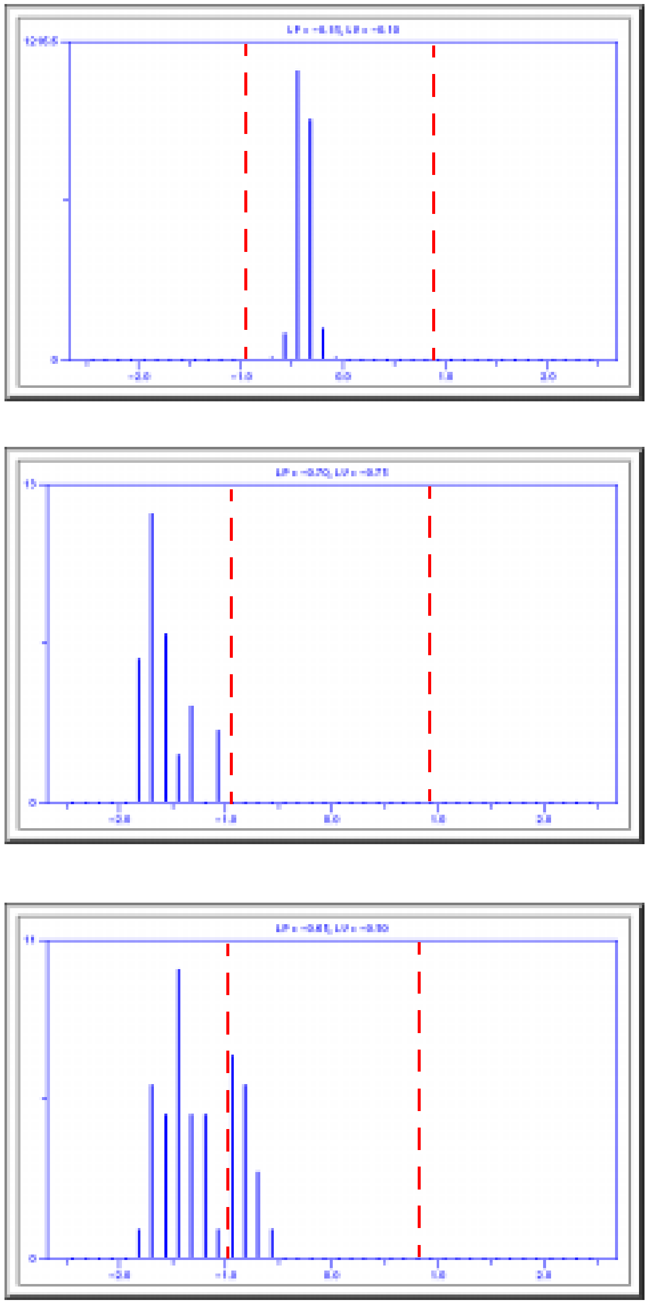

Figure 4-5: Actual future lateral position distributions for D(9) for three cases: Normal driving (top), lane change (middle), and "undecided" (bottom). The dashed lines indicate the physical lane boundary, for a 1.8m. vehicle and a 3.6m. lane width.

The top graphs shows the case of *Lp* = -0.35m and *Lv* = -0.10m/s. As the lookahead time is 1.0s, 1st order kinematics would predict that the vehicle would be at -0.45m. In this case, the prediction is accurate, as the peak of the distribution is near -0.45m, within measurement error. However, there are a significant number of cases where the future lateral position is either to the left or right of -0.45m, indicating that the driver was accelerating or decelerating, respectively. In this state, it is clear that (for a one second lookahead time), there is no reason to trigger a lane departure warning alarm.

The second case is a state which always leads to a lane change: *Lp* = -0.70m and *Lv* = -0.75m/s. In this case, the vehicle is already very close to the left lane boundary, and is heading towards it with a significant lateral velocity. As a result, the actual future lateral position is always beyond the lane boundary (after one second), indicating a lane change. This indicates that it would be appropriate to trigger a lane departure warning alarm when the vehicle is in this state, as all the evidence shows that the driver actually does depart the lane in one second. A 1st order kinematic algorithm would predict that the vehicle would be 0.55m past the lane boundary, and would trigger an alarm.

In the third case, the proper action is not obvious. Here, *Lp* = -0.65m and *Lv* = -0.45m/s. An FOD algorithm based on 1st order kinematics would predict that the vehicle would be 0.20m past the lane boundary, and would trigger an alarm. However, a significant portion of the time, the driver actually countersteers, and comes back towards the lane center. If a warning system triggers in this case, it is guaranteed to be wrong a substantial percentage of the time, resulting in nuisance alarms. In the next section, I show how to calculate and visualize this uncertainty and show how it changes over time.

4.2.7 Uncertainty

In the previous section, I showed that in certain areas of the state space, the actual future lateral position distributions make the decision to trigger or not trigger an alarm difficult. Here, I use those distributions to present a measure of uncertainty and show how it changes with FOD lookahead time.

4.2.7.1 Alarm Probability Calculation

Consider a memory table, *S*, which is trained on a particular driver and contains a set of bins, or states. Each state in *S*, s_i, contains a distribution of actual future lateral positions, as shown in Section 4.2.6. The makeup of this distribution depends on both s_i and the FOD lookahead time used in training. For the state s_i, (where s_i is a vector of lateral position and lateral velocity), it is possible to calculate $P(A_T \mid s_i)$, which is the probability of triggering a true alarm, A_T, given that the vehicle is in state s_i. Here, a true alarm is defined as a correct prediction - i.e., the future lateral position actually *does* lie beyond the virtual lane boundary. $P(A_T \mid s_i)$ is computed by calculating the percentage of the lateral position distribution which lies beyond the virtual lane boundary. This is simply:

$$P(A_T|s_i) = \frac{N(s_i)_{TA}}{N(s_i)} \tag{4-1}$$

where $N(s_i)_{TA}$ is the number of samples in the distribution which lie beyond the virtual lane boundary, and $N(s_i)$ is the total number of samples in state *s*. Given that an alarm has triggered, the probability of it being false is:

$$P(A_F|s_i) = 1 - P(A_T|s_i) \tag{4-2}$$

Where a false alarm is an alarm which is due to noisy sensor data or modelling error. Here, a modelling error is when the 1st order FOD prediction model predicts the vehicle will be beyond the virtual lane boundary, but the actual future lateral position is inside the virtual lane boundary. Using the above equation, $P(A_F)$, which is the probability of a false alarm given that an alarm event, *K*, has occurred, can be calculated as:

$$P(A_N) = \sum_{i=1}^{n} P(A_F|s_i)P(s_i|K) \tag{4-3}$$

Finally, the probability of being in a state which a 1st order kinematic vehicle state predictor triggers an alarm is:

$$P(K) = \frac{N(K)}{N(S)} \tag{4-4}$$

where *N(K)* is combined mass of the distributions which lie in states where an alarm would trigger, and *N(S)* is the combined mass of all the distributions in *S*. The next section shows how basic information theory can be used to measure the amount of uncertainty in *S*, and in following sections, I relate these metrics to alarm system performance for two of the test drivers.

4.2.7.2 Entropy Calculation

If we consider the state $s_{i,}$ which results in an alarm, as having two possible outputs, A_T or A_F, then it is possible to compute the entropy of s_i, which is a measure of the uncertainty (or, inversely, the amount of information) in s_i, as:

$$H(s_i) = -\sum_{i=1}^{2} p_i lg(p_i) \tag{4-5}$$

where $p_1 = P(A_T \mid s_i)$ and $p_2 = P(A_F \mid s_i)$, and $H(s_i)$ is the entropy. If both outputs are equally likely, then the entropy is 1, indicating maximum uncertainty, or minimum information. Conversely, if one output has a probability of 1, and the other has a probability of 0, then the entropy is 0, indicating minimum uncertainty, or maximum information.

It is also possible to calculate entropy for all the states where an alarm would trigger. This gives a measure of uncertainty for a 1st order kinematic prediction in all the cases where an alarm would trigger. I.e., if in state s_i, an alarm triggers, then $H(s_i)$ is a measure of the uncertainty of the decision to trigger. If a large majority of the actual future lateral position distribution of state s_i lies beyond the virtual lane boundary, then the decision to trigger is correct with high certainty. Similarly, if a large majority of the distribution lies within the lane boundary, then the decision to trigger is incorrect, with high certainty. The latter case is rare. What is more common is a decision to trigger when the distribution is relatively evenly distributed within and beyond the virtual lane boundary. Computing $H(s_i)$ for all *i* where an alarm triggers provides a measure of the uncertainty of the decision to trigger an alarm. This value, $H(S \mid K)$, can be computed as:

$$H(S|K) = \sum_{i=1}^{n} H(s_i)P(s_i|K) \tag{4-6}$$

where $P(s_i \mid K)$ is the probability of being in state s_i, given that an alarm event, K, has occurred. It is computed as:

$$P(s_i|K) = \frac{N(s_i)}{N(K)} \tag{4-7}$$

where $N(s_i)$ is the mass of the distribution in state s_i and $N(K)$ is the total mass of all the distributions in all states where an alarm triggered.

In the next two sections, I apply these formulations to the naturalistic data of Drivers 1 and 9, the two drivers with the most data. I will show an information theoretic basis for the differences in FOD alarm system performance between these two drivers, which was presented in Section 3.7.1.

4.2.7.3 Within Driver Results

In this section, I present two results. The first shows how the conditional false alarm probability and entropy of *S* changes as lookahead time changes (for D(1)), and the second shows the differences between applying the AutoTrak FOD parameters to D(1), vs. applying trained parameters to D(1).

Figure 4-6 shows entropy plots of *S* for D(1) over four different FOD lookahead times, 0.5s, 1.0s, 1.5s, and 2.0s. The virtual lane boundary used in all cases was 0.1m, which is what AutoTrak uses. The lightest gray area in each plot indicates states which have data within them. The darker areas indicate areas of higher entropy. As the lookahead time increases, the areas of higher entropy grow, indicating increasing uncertainty.

The central areas have low entropy, because in most cases, when the vehicle is in one of those states, there is no impending lane change. Similarly, the bottom left and top right portions of the state space, which correspond to lane changes, also have low entropy. In these states, it is unlikely that drivers will correct their steering. Rather, they are most likely in a lane change, and will continue along such a trajectory.

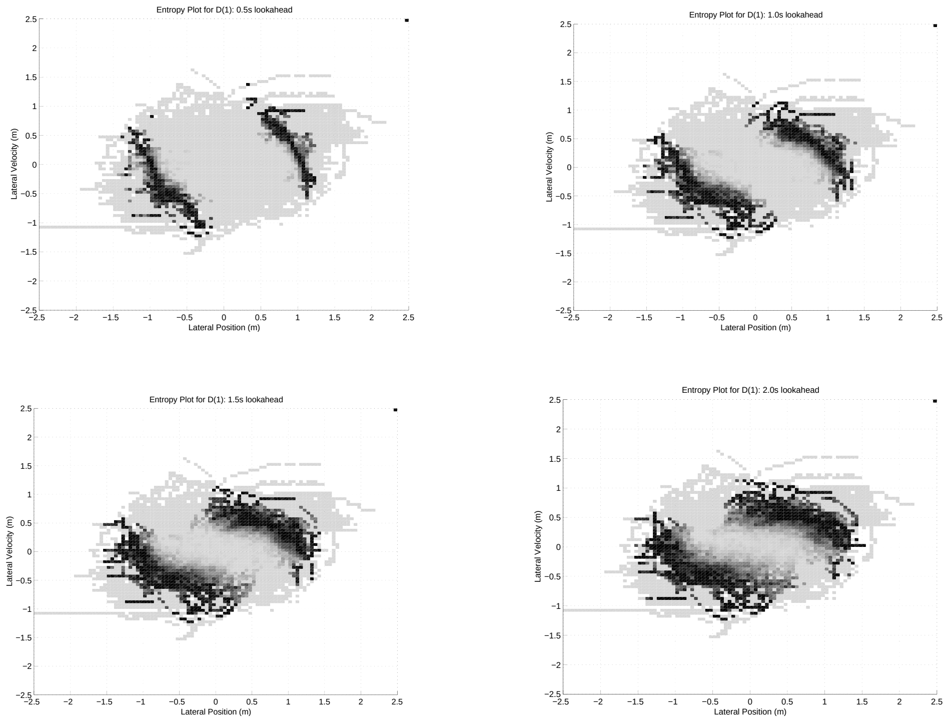

Figure 4-6: Entropy plots for D(1). The lookahead time is 0.5s, 1.0s, 1.5s, and 2.0s, respectively. Dark areas indicates states with high entropy.

The areas with high entropy lie along the boundary between alarm trigger and no trigger. In these areas, it is difficult to predict the future action of drivers, as a given trajectory may either be part of a lane change, or just normal swerving. Therefore, the uncertainty is higher, and attempts to trigger an alarm in these states lead to false and nuisance alarms.

Table 4-1 shows $P(K)$, $P(A_F)$, and $H(S \mid K)$ for various lookahead times for D(1). As lookahead time increases, so does P(K), which is the probability of being in a state which triggers an alarm. This matches the observation in Figure 4-4, where the trigger areas grew with lookahead time. With a 0.10m virtual lane boundary, a 3.0s lookahead time means that the driver would be triggering an alarm 17% of the time. $P(A_F)$, the probability of a false alarm, also increases with lookahead time. This indicates that predictions become less and less reliable as we try to predict further into the future. For a 0.10m virtual lane boundary and 1.0s lookahead time, the odds are almost even as to whether a trigger is appropriate.

FOD Lookahead Time (s)	**P(K)**	**$P(A_F)$**	**H(S | K)**
0.0	0.02	0.00	0.00
0.5	0.03	0.26	0.51
1.0	0.05	0.47	0.64
1.5	0.07	0.60	0.68
2.0	0.10	0.72	0.62
2.5	0.13	0.77	0.56
3.0	0.17	0.82	0.50

TABLE 4-1 : Alarm Probability, False Alarm Probability, and Entropy Calculations for D(1)

H(S | K), which is the weighted sum of the entropy over all the trigger states, starts to increase with lookahead time, but then begins to decrease. The reason for the decrease is the increase in *P(K)*. As the lookahead time increases, an alarm triggers in more states. In most cases, the driver doesn't maintain a constant linear velocity over the lookahead time, but rather corrects and returns towards the lane center. Therefore, as lookahead time increases, the likelihood of an incorrect trigger increases as well. After a certain point, alarms start triggering in states where the driver isn't even remotely likely to actually exceed the virtual lane boundary. Therefore, the entropy in these states is actually lower; i.e, we are nearly sure that we are making a wrong decision.

Figure 4-7 shows kinematic sensitivity and entropy plots for two difference cases: 1) the application of the AutoTrak FOD parameters to D(1), and 2) the application of one of the trained FOD parameters to D(1). The trained FOD parameters are lookahead time = 2.0s and virtual lane boundary = 0.9m. These parameters were picked from Figure 3-21, and provide the same warning onset time as the AutoTrak parameters, with a lower nuisance alarm rate. Of the trained parameters in Figure 3-21, the parameters I picked are furthest away from AutoTrak. I did this to emphasize the difference between the trained parameters and the AutoTrak parameters. Why is the nuisance alarm rate lower? In Section 3.7, I claimed that it was because the trained parameters allowed the driver to swerve beyond the lane boundary

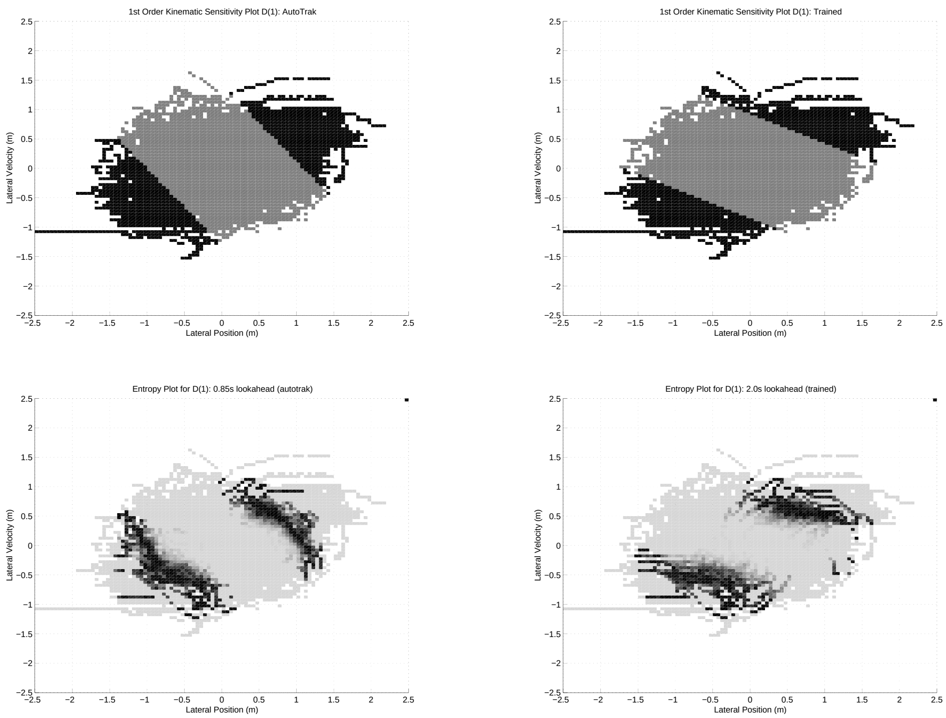

Figure 4-7: Alarm Trigger Sensitivity and Entropy plots for AutoTrak and Trained FOD parameters, applied to D(1).

Parameter	P(K)	P(A_F)	H(S | K)
AutoTrak	0.04	0.36	0.59
Trained	0.01	0.35	0.53

TABLE 4-2 AutoTrak vs. Trained probabilities and entropy for D(1)

more than the AutoTrak parameters did. Figure 4-7 and Table 4-2 show that this is true. The table shows the alarm probabilities, conditional false alarm probabilities, and entropy of the AutoTrak and trained parameters on D(1). While the trained parameters have a slight reduction in false alarm probability and entropy, the big difference is in *P(K)*, the probability of being in an alarm state. For the trained parameters, *P(K)* is much lower. This means that the trained FOD parameters are overall less likely to cause a trigger than the AutoTrak parameters. This is because the trained parameters rotate the decision boundary to avoid areas of high lateral position and low lateral velocity, allowing D(1) to swerve more. Therefore, nuisance alarms are less likely. However, for a range of lateral velocities around 0.7m/s, the trigger point (and corresponding warning onset time) is the same as AutoTrak.

This reduction in nuisance alarms comes at a cost of later detection of slow drifts. From the sensitivity figures, we can see that the vehicle has to drift further away from the lane center before an alarm triggers in low lateral velocity cases. For instance, for a lateral velocity of 0.20m/s, AutoTrak would trigger when the vehicle is about 0.90m away from the lane center. However, the trained parameters would trigger at around 1.3m. In the opposite direction, the trained parameters would trigger *earlier* than AutoTrak for high lateral velocity. What is happening is that the trained parameters pivot the boundary line about a point. This point is equal to the average lateral velocity during lane changes. The direction of the pivot (towards decreasing slope) maintains performance in the region of lane change lateral velocity, but changes performance outside that region. The lower slope of the trained parameters actually increases the variance of the trigger point, but, around the region of lateral velocities which encompass most lane changes, the trigger point does not vary much. This increased variance provides the benefit of reduced nuisance alarms (alarms which are not lane changes): 106 for AutoTrak, vs. 45 for the trained parameters (when using an alarm decision model which does not account for curve cutting or local adaptation). This is a better than two-fold reduction in nuisance alarms, and is due to the lower probability of being in an alarm state.

4.2.7.4 Between Driver Results

The probability and entropy measures described above can also be used to investigate differences in FOD alarm system performance between drivers. In particular, I concentrate on D(1) and D(9). There is more data available on these two drivers than on any of the others, and their warning system performance is quite different. Using the full alarm decision model, applying the AutoTrak parameters to D(1) results in 37 nuisance alarms over nearly 5.5 hours, whereas applying the same parameters to D(9) results in only 6 nuisance alarms over nearly 7 hours. However, the warning onset times are similar.

Driver	P(K)	$P(A_F)$	H(S | K)
D(1)	0.04	0.36	0.60
D(9)	0.02	0.12	0.24

TABLE 4-3 : Entropy and Probabilities for D(1) and D(9), using AutoTrak FOD Parameters.

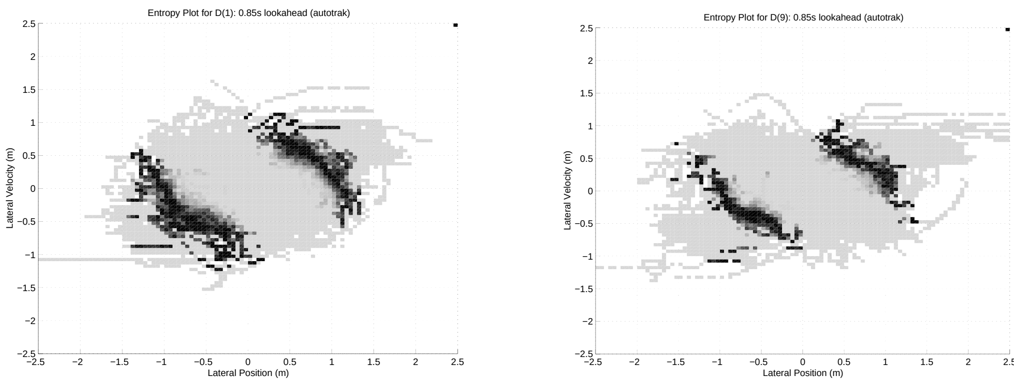

Figure 4-8: Entropy Plots for D(1) and D(9). Dark areas are areas of high entropy

Why is this? Figure 4-8 and Table 4-3 provides some insight into the reason for D(9)'s improved performance. Figure 4-8 shows entropy plots for D(1) and D(9), using AutoTrak FOD parameters. There are fewer areas of high entropy for D(9), as compared to D(1). This means that in general, kinematic predictions for D(9) tend to have higher certainty. This higher certainty could indicate a higher certainty of being *wrong*. However, Table 4-3 shows that this is not the case. The conditional probability of a false alarm, $P(A_F)$, is one-third that of D(1). This indicates that when an alarm triggers, it is much less likely to be a false (and therefore nuisance) alarm for D(9), than for D(1). This is reflected in the entropy measures as well, where D(9)'s entropy is significantly lower than D(1).

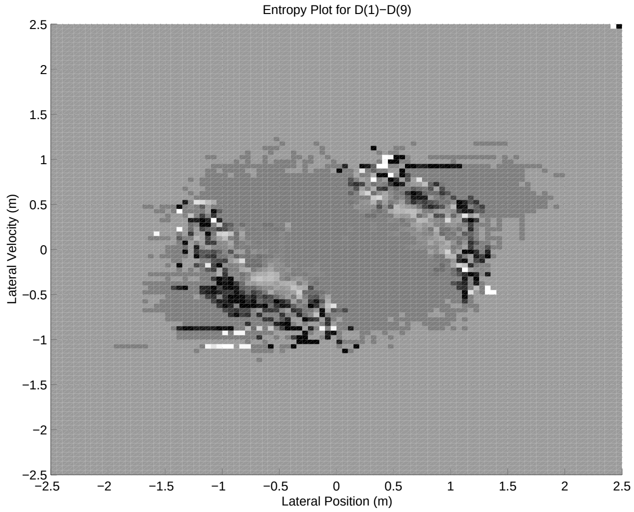

Figure 4-9: Entropy plot of D(1)-D(9). Dark areas are where D(1)'s entropy is greater, white areas, where D(9) is greater.

D(9) is also less likely to be in a trigger state than D(1). This is because in general D(9) doesn't weave on the road as often or as excessively. Therefore, this subject doesn't enter alarm states as often as D(1), and when he does, it is most likely while changing lanes. D(1), in contrast, is in states which trigger alarms much more often, even when not changing lanes. Therefore, D(1) has a higher nuisance alarm rate.

Figure 4-9 shows an entropy plot which corresponds to subtracting the entropy at each state, s_i, of D(9) from the entropy of the corresponding state in D(1). Dark regions indicate areas where D(1)'s entropy is greater than D(9), and light areas indicate the opposite. The comparison was only done on states which had data in both D(1) and D(9). Overall, there are more dark areas than there are light, and the magnitude of the difference indicates that D(1) has higher entropy than D(9), in a larger number of states. The average difference over all states where there is any difference in entropy is 0.18, indicating that on average, D(1)'s entropy is higher.

Left lane changes and deviations seem to be an area where D(1) and D(9) greatly differ. D(1)'s entropy around the area L_p = -1.0m and L_v = -0.5m/s is much higher than D(9)'s. When near that state, an alarm is almost always appropriate for D(9), as seen by D(9)'s low entropy in that area. D(1), however, tends to be in that state even during normal driving, which increases the uncertainty and nuisance alarms.

4.3 Summary

In the above sections, I introduce a new method, derived from memory based learning, to visualize and analyze driver behavior. This allows me to show that in certain areas of vehicle state space, the likely future lateral position of the vehicle cannot be predicted with great accuracy. I use this, along with calculations of false alarm rate and entropy, to examine two questions: 1) Why does training FOD parameters help Driver 1's performance, and 2) Why does Driver 1 have many more nuisance alarms than Driver 9?.

The answer to the first question involves looking at the states in which an alarm would trigger, for both the AutoTrak FOD parameters, and the trained parameters. As a 1^{st} order vehicle state predictor is used, the boundary, in vehicle state space, between trigger states and non trigger states is linear. The slope and location of this boundary is determined by the FOD parameters. The trained parameters end up pivoting the boundary line about a point near 0.7m/s, which is the average lateral velocity during D(1)'s lane changes. The result of this pivoting is to change previous trigger states with high entropy (and thus high nuisance alarms) into non-trigger states, which reduces the overall nuisance alarm rate. This, however, comes at the expense of trigger point variance. With the trained parameters, the location (relative to lane center) of alarm triggers changes with lateral velocity *more* than it does with AutoTrak.

This trade-off suggests an area for future human factors research. Namely, what is more important, alarm system predictability, or low nuisance alarm rate? For most drivers, Table 3-7 shows that with an enhanced alarm decision model which accounts for curve cutting and local adaptation behaviors, the AutoTrak parameters provide a low nuisance alarm rate, while maintaining predictability. For certain types of drivers, who weave excessively, training FOD parameters provides a benefit, at the expense of predictability. The answer to the above question is beyond the scope of this thesis, but the question is still an interesting one.

I also use the tools I develop to explore differences in behavior between two drivers who have very different alarm system performance. D(1) has many nuisance alarms, while D(9) has few. There are two main differences: nuisance alarm probability, and entropy. These two differences are related. D(9) has a lower nuisance alarm probability than D(1). I.e., given that an alarm triggers, the alarm is much less likely to be a nuisance alarm for D(9) than it is for D(1). This leads to a difference in entropy levels as well. D(1) enters states that normally trigger an alarm even during the course of normal driving. This increases the uncertainty of the trigger. I then compare the entropy of D(1) and D(9), and show that both drivers differ when making left lane changes. D(1) seems to weave a lot, making those lane changes hard to accurately predict.

This prediction is mainly difficult due to the limited state information which is available. Lateral position and lateral velocity do not provide enough information to accurately predict future vehicle state over long periods of time. This is because drivers change their steering behavior over longer lookahead times, violating the constant lateral velocity assumption. One possible solution is to use additional state information, such as lateral acceleration. Lateral acceleration could provide information about when the driver’s lateral velocity is increasing or decreasing, which could help filter out alarm triggers that occur during normal driver behavior.

Another possibility is to use Fuzzy Logic methods [34] to account for variability in driver behavior. It would be possible to conceive of classifying drivers into categories such as “very tight,” “tight,” “loose,” or “dangerous” based on membership sets representing those categories. Lateral position and lateral velocity could be used to create the membership sets based on training data. Alternatively, the cells in the MBL table represent quite small areas of state space. A fuzzy combination of nearby states could allow one to reduce the granularity of the state table, thereby reducing noise, while maintaining accurate representation of likely future position.

CHAPTER 5 Conclusion

5.1 Thesis Summary and Contributions

The major contribution of this thesis is the development and evaluation of a warning algorithm with a novel alarm decision model, trained to a particular driver, and evaluated on real world driving data, as mentioned in my thesis statement in Section 1.1. This decision model determines when the driver is in danger of a run-off-road (ROR) crash, based on the use of lane changes as simulated substantial lane departures. This model is then used in the Future Offset Distance (FOD) warning algorithm, along with a linear vehicle state predictor, to reduce nuisance alarms while maintaining adequate warning time.

The vast majority of the previous work has concentrated on data collected on driving simulators. While simulators have advantages, their fidelity is often questionable, which makes interpreting results difficult. My use of real world data means that the results presented in this thesis can be expected to carry over to on-line lane departure warning systems (LDWS). Additional contributions include a methodology for the collection of controlled and naturalistic real world driving data, and a low-level analysis of the performance of a LDWS on real world data.

In chapter 2, I describe two data collection efforts, which resulted in a combined 100 hours of driver data. About 50 hours of this data is high quality and usable. Of the 50 hours, 30 were collected under controlled circumstances. The route, time of day, vehicle type, and experimenter presence were all the same. As this data was collected on NavLab 8, a sensorized minivan, there is also information on surrounding traffic, which was obtained by radar sensors, along with information on the driver's steering input. I also collected 30 hours of video of the driver's view, which would be valuable to future researchers who wish to make use of this data.

The other 20 hours of high quality data are of drivers behaving naturally, with no constraints on vehicle route, road conditions, or experimenter presence. While the Carnegie Mellon Research Institute has collected long-term data on truck drivers, no other study has collected as much unconstrained data on passenger car drivers. This data has been highly valuable in the development and evaluation of my individually trained warning algorithm.

This warning algorithm, along with its associated alarm decision model, is the focus of Chapter 3. The previous work in this area has either been too simple to capture realistic driver behavior, such as the Roadside Rumble Strips and Time to Lane Crossing (TLC), or complicated enough that their ability to work with realistically noisy data is unproven and questionable. These systems, which I describe at the beginning of the chapter, suffer from either low warning onset time, or high nuisance alarm rate. I determine this by applying the roadside rumble strips and TLC algorithms to the naturalistic data described in Chapter 2. For this evaluation, and all the others, I substitute lane changes for true dangerous deviations for three reasons: 1) the data did not contain any actual crashes, 2) lane changes are similar to run-off-road incidents caused by unintended steering wheel movement, and 3) the data contains over 600 lane changes.

I then introduce a new warning algorithm, Future Offset Distance (FOD), which adds a virtual lane boundary parameter, allowing drivers to weave more aggressively without triggering an alarm. This is followed by a description of two algorithms to train the FOD model. The first trains the FOD model on a collection of drivers, and the second trains the model on an individual driver. This is done to determine whether training for an individual driver provides any benefit over training on an "average" driver, which is a combination of drivers. The results of these experiments are in terms of Warning Onset Time (WOT) and Nuisance Alarm Rate (NAR), which are defined in Section 1.2. Intuitively, WOT is the amount of time drivers would have to react to a warning, before their lateral position exceeds the physical lane boundary by a pre-defined (currently 3ft) threshold. These results are compared against the hand-tuned FOD parameters which Pomerleau's AutoTrak LDWS uses. The results show a significant improvement in NAR with a minimal drop in WOT for one of the five drivers, when using individually trained parameters, and comparable performance on the other four.

The FOD parameters can only be trained to find an average best fit for a particular driver. It does not account for local deviations in driver behavior, due to road geometry, the presence of other vehicles or obstacles, and changes in driver behavior over time. Therefore, I model two additional driver behaviors which are often anecdotally observed, but until now, have not been quantified. These are curve cutting, and local deviations.

The first, curve cutting, is the tendency for drivers to shift towards the inside of a curve. This increases the overall radius of curvature of the vehicle, and reduces the lateral forces, increasing occupant comfort. I use the naturalistic data to show that drivers have a tendency to cut towards the inside of a curve, and that the amount of curve-cutting varies inversely with the radius of curvature. I use this evidence to add a term to the alarm decision model to account for curve cutting, and find appropriate parameters for the term by exploring the parameter space. Modelling curve cutting, it turns out, has an overall beneficial impact on warning system results, by reducing the number of nuisance alarms without adversely affecting the warning onset time. The warning time *could* be reduced if drivers drift off the inside of a curve. However, such accidents are four times less likely than accidents due to drifting off the outside of a curve.

The second driver behavior I add to the alarm decision model accounts for short term changes in the vehicle's mean lateral position. These short term shifts can be caused by many reasons, including the presence of adjacent vehicles, construction zones, or changes in the driver's state. The naturalistic data does not contain information on any of these factors, although inspection of the data reveals stretches where drivers temporarily shift their mean lateral position. Appendix C provides some evidence for this phenomenon in the presence of adjacent vehicles using the NavLab 8 data. To account for short term variations, I add in a term to increase the virtual lane boundary on the side to which the vehicle has shifted. This has the effect of giving the driver more leeway on that side. It avoids situations where a driver in the left lane may be passing a truck, which causes the driver to shift to the left, which triggers an alarm, and sends the driver back to the right, approaching the truck, and possibly causing a crash. Running the LDWS experiments on the stored datasets with this complete alarm decision model further lowers the nuisance alarm rate, while sparing the warning onset time.

The overall results of Chapter 3 show that improving the alarm decision model causes an overall improvement in Future Offset Distance LDWS performance. Individually training this model helps for one of the five drivers, D(1), while its effects on the other four drivers are smaller. Part of the reason for this is that the other drivers do not have nearly as many nuisance alarms as D(1), so there is not as much room for improvement. Moreover, there are only two drivers for whom I had more than 5 hours of data, D(1) and D(9), and D(9) has a very low nuisance alarm rate in nearly all cases. Performance of the individualized model for other drivers may have been better had there been more data available.

Chapter 3 raises two questions: 1) Why does an individualized FOD model improve LDWS performance for D(1), but not other drivers, and 2) Why does D(1) have so many more nuisance alarms than other drivers? Chapter 4 answers these questions, and brings forth interesting insight into the behavior of linear warning algorithms and alarm decision models.

To answer these questions, I use a memory-based learning (MBL) framework, which allows me to look at driver behavior at a much finer scale than others have done in the past. To do this, I discretize the vehicle state space (lateral position and lateral velocity), and create bins for each state vector. These bins contain the *actual* future lateral position of a vehicle which is in a state corresponding to the bin’s state vector, after a specified lookahead time. After adding hours of data into this memory table, each bin contains a distribution of actual future lateral position. If this distribution contained a single value, that would be evidence that future vehicle state is completely predictable. However, the distributions generally have a gaussian shape, indicating that there is uncertainty in predicting future driver position. In certain areas of the state space, the distributions are bimodal, due to the driver either making a “corrective” action to avoid a road departure, or continuing with a lane change. This bimodality makes accurately predicting future vehicle state very difficult in certain cases.

I use an FOD sensitivity analysis, on top of the MBL framework, to examine why a trained model improves performance for D(1), but not for D(9). This analysis demonstrates which areas of the vehicle state space would cause a trigger given a particular set of FOD parameters. The trained parameters, it turns out, allow for more swerving than that AutoTrak parameters, by not triggering as early when the vehicle crosses the lane boundary at low lateral velocities. The trained parameter's response to higher lateral velocity deviations, such as during lane changes, is mostly unaffected, resulting in an equivalent warning onset time.

D(1)'s excessive swerving has another effect on LDWS performance, which is not immediately obvious. This effect is higher uncertainty in future state prediction, relative to other drivers. To quantify this, I use an entropy metric, which is a measure of the amount of information or uncertainty in a distribution. I calculate an entropy measure for each vehicle state for D(1) and D(9), measuring the uncertainty of a decision to trigger an FOD alarm, based on the probability of the *actual* future lateral position being beyond the virtual lane boundary (which, remember, is the criteria for triggering an FOD alarm). This analysis shows that D(1) has an overall higher uncertainty than D(9). This means that the decision to trigger an alarm is wrong more often for D(1) then for D(9), which leads to more nuisance alarms. This presents a fundamental limit to the performance which can be achieved for D(1).

The main conclusions which can be drawn from this thesis are: 1) Real world data is important for the evaluation of lane departure warning systems. 2) State of the art warning systems, such as AutoTrak, perform pretty well. However, it is possible to improve performance for drivers who weave a lot through individual training and an improved alarm decision model. 3) The use of a simplified vehicle state space, namely lateral position and lateral velocity, while convenient in terms of lane tracker design, may not provide enough information to help all drivers. Finally, 4) A great deal of additional driver data is needed to defend the statistical significance of the results presented here. These conclusions lead me to recommendations for future work, which is presented next.

5.2 Future Work

Here, I present directions for future work in two areas: Data collection and field evaluation studies, and warning system design.

5.2.1 Data Collection and Evaluation

5.2.1.1 More Data

The lack of data is a limiting factor in evaluating the performance of trained lane departure warning algorithms. While two of the five available drivers had over five hours of data, the other three drivers only had 2-3 hours of data. This means that there may not have been enough training data to learn FOD models for those drivers, and that medium scale variations (on the order of a half hour or so) in behavior could have a significant impact on performance. With more training data, these variations would not have as great an influence on the results. For example, consider the leave-one-out training algorithm described in Section 3.5.4. If there are two hours of data available for a subject, then for any given half hour test set, the training set is 1.5 hours. If a particular half hour test set has a significantly different mean lateral position or lateral position variance than the training set (due to road conditions or weather, perhaps), then the trained parameters could perform poorly on the test set. With significantly more data (5+ hours), this effect is not as strong as with only two hours of data.

Ideally, it would be nice to have 8+ hours of data per driver, over the course of a week or more. This would allow investigations into how driving behavior changes over a larger time scale, along with looking at how any observed long term trends vary between drivers. Similarly, it would provide a firmer answer for the question of whether training once is sufficient for a particular driver, or whether a longer term adaptation of FOD parameters is necessary. While the naturalistic data was collected over multiple days of driving, the data collection system could not timestamp the data, so it was impossible to determine when the data was collected. Moreover, the lack of data means that at best, only 3 hours or so of data was collected per day. Having more data would be a benefit in assessing the statistical significance of the results presented in Chapter 3.

5.2.1.2 More Subjects and Larger Demographics

Although the naturalistic data collection effort resulted in 70 hours of data from nine drivers, a majority of the data turned out to be unusable, due to a combination of system errors and low-confidence data. In the end, only five drivers provided any usable data at all. These five drivers are not enough, as mentioned above, to determine statistical significance. However, the experimental design I put forth in Section 3.4, applied to long term (1-2 weeks) data on a larger number of drivers could provide a definitive answer to what percentage of drivers could be helped by an individually trained FOD model, and what characteristics such drivers need to possess.

All nine drivers who volunteered for the naturalistic study, and 18 of the 20 NavLab 8 subjects are members of the CMU Robotics Institute. The two subjects in the NavLab 8 study who are not members of the Robotics Institute are still employees of CMU. This is an insurance requirement, so recruiting subjects from outside of CMU is not possible. Therefore, the demographics of the subjects are fairly limited. All were professionals in the 24-45 age group. There are no teenagers or senior citizens in the subject pool. These two groups tend to have higher accident rates; teenagers because of inexperience, and senior citizens due to age-related issues such as poor eyesight and reduced reaction time. It is likely that including subjects from those groups would reveal greater variance in driver behavior. Similarly, including people from non-academic backgrounds, especially people who do not often drive on long trips, would provide interesting insight into driver behavior.

5.2.1.3 More State Information

AutoTrak makes three independent vision-based measurements: lateral position, lane width, and boundary type. This information is transformed to provide other state variables such as lateral velocity, and vehicle relative heading. Longitudinal velocity is available when a GPS option is installed, but not all the systems used in the naturalistic study had this option available.

Having more state information would increase the accuracy of the FOD state predictor. In Chapter 4, I show how the actual future lateral position for a particular state can be bimodal, where each mode reflects the driver's choice to depart the lane (in a lane change), or correct and re-center the vehicle. This bimodality presents a limitation on the certainty of any warning algorithm which only uses instantaneous lateral position and lateral velocity information. A system which had access to lateral acceleration would be able to distinguish between the two cases earlier, and only trigger an alarm when the driver was actually departing the lane. This is because when the driver is correcting, acceleration should be less than zero, whereas when the driver is departing the lane, acceleration should be close to zero (for slow drifts or constant speed lane changes), or increasing.

Generating reliable lateral acceleration information is, however, quite difficult. The main problem is that lateral acceleration needs to be computed relative to the lane edge. Accelerometers would provide acceleration information along one of the vehicle axes, and this axis may not be perpendicular to the lane edge. This is because accelerometers measure acceleration along an inertial frame, whereas what is required is acceleration in the road's coordinate frame. That is, if the driver were perfectly tracking the road, and following a curve, then an accelerometer would still report acceleration, even though the actual acceleration relative to the lane edge is zero. This means that very accurate registration is required, to compute the yaw of the vehicle relative to the road, and the upcoming road curvature. While AutoTrak does provide heading information, it is based on the lateral position measurements, so any noise in this measurement would carry over to the heading, and would be compounded by the noisy outputs of most accelerometers. There is a further problem with the use of accelerometers, which is caused by superelevation. Superelevation is the bank in a curve, which is used to minimize slipping. This banking changes the acceleration measurements, as now a fraction of g is introduced. It would be possible to factor this out, except that the amount of superelevation between curves is generally not constant, and is usually undocumented.

An alternative is to differentiate the lateral velocity output. This would provide acceleration information in the appropriate coordinate frame. However, the lateral velocity output is the derivative of the lateral position measurements. Differentiating the lateral position estimates twice introduces a great deal of noise, making the signal useless.

AutoTrak does provide other information, which, in an online system, can be used to suppress nuisance alarms. AutoTrak can distinguish between solid and dashed lane boundaries. Therefore, warnings could be suppressed if the vehicle were crossing a dashed lane boundary, indicating that it was changing lanes. Similarly, it is possible in the current version of AutoTrak to detect the state of a vehicle's turn signal, and suppress alarms if drivers are indicating that they are changing lanes. I did not use lane boundary information because I wanted to trigger on lane changes, as they were my surrogate true alarms. In an online system, however, suppressing alarms (or altering their modality) during lane changes is useful.

Another type of state information which would be very useful would be information on adjacent vehicles. While this is impossible to get with forward looking vision systems, other sensors, such as radar, can be used to detect other vehicles surrounding the driver. This is because drivers tend to change their lane keeping behavior when passing other vehicles, or when there is another vehicle along side them. I attempt to account for this in Section 3.6.2 with a local disturbance model, which adapts the virtual lane boundary based on an updated estimate of the driver's mean lateral position. However, having information about surrounding vehicles, or the presence of construction zones would be quite useful. While adding radar sensors to vehicles can be expensive for production model vehicles, current high end cars are starting to include an adaptive cruise control option, which makes use of a front mounted radar to slow the vehicle down to prevent collisions. Similarly, new Ford Windstars have a rear collision warning option, for short range obstacles such as when parallel parking. It is conceivable that given a proven benefit, auto manufacturers may someday include side sensing radars. These radars only need to be able to detect the presence of another vehicle, and therefore would be cheaper than adaptive cruise control radar. They could also serve double duty as a blind spot warning sensor, to prevent drivers from changing lanes or merging into overtaking vehicles.

5.2.1.4 More Road Information

AutoTrak provides lane width and road curvature information, which is useful. However, other information, such as the lack or presence of a shoulder, and shoulder width would be nice to have. For my experiments, I assume a shoulder width of three feet. In many cases, the actual shoulder width is greater, particularly adjacent to the right lane. A wider shoulder can influence driving behavior. There is evidence [57] that when driving on rural roads, which usually have a very small shoulder, or no shoulder at all, people tend to drive more carefully, and with a lower lateral position variance. It is possible, although not proven, that when on highways with large (say greater than 1.5 meters) shoulders, drivers may drive more loosely, as there is little danger of drifting off the road. My experiments do not account for that effect because shoulder width information is not available. It is possible that such information will be available in the future, as GPS-registered road maps are becoming more popular and vision systems are becoming more capable. Current road databases do not contain shoulder width information, but if and when that changes, taking advantage of that information would provide a benefit.

5.2.1.5 Image Data

During the NavLab 8 study, I collected video of the road ahead during all the trials. During later analysis, this was useful to have, as I could look through the data, find interesting events, then look at the tape to see what was actually going on. I could not collect video information on the naturalistic study subjects. It would be useful to be able to log image data, either as video, or as still pictures, particularly in “interesting” cases, such as when the AutoTrak alarm triggers. It is now possible to collect low resolution digital images onto flash card with AutoTrak, but at the time of the naturalistic study, this capability did not exist. If images could be captured every time the system believes the driver is near to crossing or has crossed the physical lane boundary, then the areas where most nuisance alarms occur would be captured, and could be looked at in further analysis, to try to determine whether there were any extenuating circumstances for the trigger.

5.2.1.6 Lane Changes as Surrogate Lane Departures

One difficulty in trying to prevent run-off-road (ROR) crashes is that they are quite rare. Therefore, finding positive examples of ROR crashes is difficult, necessitating the use of alternative events to simulate the crashes. I chose to use lane changes, due to their similarity to certain types of ROR crashes, and the large number of lane changes which most drivers make during a long stretch of driving. Lane changes, however, are *not* ROR crashes, and it likely that there are factors in vehicle trajectory and driver behavior (such as drowsiness) during ROR events which do not exist during lane changes.

Future work will have to address this issue, by attempting to capture true ROR events. In the 110 hours of data I have collected, there is not a single ROR crash. To have any hope of collecting data on true crashes, it is likely that data collection efforts have to increase by several orders of magnitude or more. Having a great deal more data would also increase the number of significant departures, which would be useful to have. It would also be useful to have video of the drivers' faces, to see what their state is during interesting events, or perhaps steering wheel input, to determine whether they made a steering correction. This could help determine whether an alarm is appropriate in a particular situation.

5.2.2 Warning System Design

5.2.2.1 Explicit Dynamic Models

In this thesis, I have not used any explicit modelling of the driver, vehicle, or road. I use a purely data-centric approach. However, it is possible to build actual driver models and vehicle models, fit them to actual data, and use forms of parameter analysis to determine whether differences in model fit are statistically significant. Pillutti [58] has done this for simulated data. One advantage of this is that is could perhaps work with less training data (assuming drivers can be modelled as stationary processes), and could be used to develop models of driver and vehicle interaction explicitly, rather than implicitly, as I do now.

One of the difficulties in this approach, however, is the effect of road conditions on vehicle model performance. It is not enough to have just a good vehicle model. Rather, it is critical to have data on how the vehicle interacts with the road surface, and this requires weather information on rain or snow precipitation, or the presence of ice on the road, or even the type of material used to construct the road.

5.2.2.2 Reaction to Warning

In this thesis, although I make references to human factors issues, I do not address them directly. One of the most important factors is how a driver would react to an alarm. There are many possibilities. The "unsafe nuisance alarm" definition is an example of a negative outcome. It is possible that a loud audible alarm could cause drivers to overreact, and place themselves in dangerous situations. This is obviously an undesirable outcome. There has been unpublished work by the Vehicle Research Test Center (VRTC), using their DAS-CAR system, which shows that initially, drivers may freeze for up to two seconds after hear-ing an alarm from an unknown source. They spend the time looking for the source of the alarm. For a lane departure warning, freezing for two seconds could be enough time for the driver to completely depart the road, or enter a situation from which a recovery is impossible.

5.2.2.3 Warning Modality

An audible alarm is not the only way to warn a driver of an impending lane departure. Haptic warnings are another modality which could be beneficial to explore. The previous work in this area [2] [54] [72] has usually used an airplane “stick shaker” model, which nudges or vibrates the steering wheel in the direction opposite the lane edge, alerting the driver to turn in that direction. It is possible that such a method would be more intuitive to the driver, and not cause the initial confusion that the VRTC group has reported for audible alarms. The Odetics warning system, which is scheduled to be offered on Mercedes heavy trucks this fall, simulates a rumble strip sound in the direction of the departure, taking advantage of the fact that most people, especially truck drivers, have experienced rumble strips in the past, and know how to react to them. The main problem with haptic interfaces is that they require vehicle modification. An audible system such as AutoTrak has an advantage that it is easy and cheap to install, as no vehicle modifications are necessary. However, there may be advantages to haptic warnings over audible warnings, which need to be explored.

Another issue, which is related to warning modality, is warning intensity. Currently, AutoTrak provides two levels of warnings. The first and lowest level optionally triggers when the driver changes lanes. This is mainly included as a “comfort” feature, so the driver knows the system is working. The second level of alarm is louder and more intensive, and occurs when the driver deviates onto the shoulder. It is possible to imagine a different hierarchy of warnings, which trigger based on the perceived danger. The memory based learning approach provides a mechanism for doing just that. Entropy measurements of the training data could be used to determine certainty estimates of the decision to trigger a warning. When the certainty is low, then the system is unsure about the safety of the situation, but realizes that there is the potential (based on past training data) for danger. In this case, a small visual warning could be triggered, possibly on a heads-up display in front of the driver. If the deviation continued, then the certainty of danger increases, and more intrusive warnings could be triggered, culminating in a loud audible alert similar to AutoTrak’s. Such a system may have an overall effect in reducing audible alerts, as drivers could take corrective action at lower certainty levels. This could change the role of the warning system to one which is more advisory. The driver is free to ignore unobtrusive visual warnings, but only up to a point. It is also likely, although not demonstrated, that nuisance visual warnings would be less annoying that nuisance audible warnings, which could increase driver acceptance of an LDWS.

5.2.2.4 Predictability vs. Optimality

The experimental results I describe in Section 3.7, based on roadside rumble strips and TLC, have either low warning time, or high nuisance alarm rate, respectively. Similarly, for drivers who weave a great deal, the FOD algorithm can produce an unacceptable number of nuisance alarms. Based on the results of D(1), it appears likely that learning FOD parameters can improve the nuisance alarm rate. However, the “fixed parameter” methods have a possible advantage over an individually trained FOD algorithm with an advanced alarm deviation model; namely, the behavior of fixed parameter systems is highly predictable. At the extreme, there is the roadside rumble strip system. Drivers know that when their tire hits the rumble strip, a warning will sound. This is entirely predictable, and could provide a measure of comfort. Similarly, AutoTrak is tuned to sound a warning when the vehicle’s tire is just about to cross the lane boundary. If drivers are able to predict the behavior of the warning system, they may feel more comfortable with it. In the case of training FOD parameters for D(1), the sensitivity analysis in Section 4.2.7.3 shows that the trained parameters suffer from an increased variability in the “alarm trigger point”. For low lateral velocities, the driver can drift further away from the lane than for high lateral velocities. The curve cutting and local adaptation behaviors add to this variability. While these extensions help reduce nuisance alarms, they increases the unpredictably of the system. It is possible, although not demonstrated, that drivers may prefer predictability over performance. An online user study, using a predictable warning system, such as rumble strips, vs. a more complicated warning algorithm, such as the one I have presented, would provide valuable insight into the preferences of a driver.

5.2.2.5 Driver Adaptation Issues

The driver adaptation methods and new alarm decision model which improve warning system performance for D(1) could be taken to an extreme. By adapting to D(1)'s weaving, the LDWS is tacitly validating D(1)'s driving behavior. The same adaptation could be used to adapt to *really* dangerous driving behavior, reducing the value of a the LDWS, or even making it dangerous, as the driver could get a false sense of security. This leads to the question -- how much adaptation is too much? Answering this question would require a great deal of real world driver data, over broad subject demographic, vehicle type, and weather condition, to determine the normal operating envelope of drivers. When such information becomes available, then it would be possible to limit the adaptation of a warning system, so that the extremes of driver behavior are not tolerated. The results on the five naturalistic subjects show, however, that adapting to a driver, or at least providing a manual sensitivity adjustment, is a useful tool.

In a similar vein, providing constant feedback about drivers' performance may help them improve their driving. The AutoTrak system displays a constantly changing "Alertness Index," which is a function of the driver's lateral position variance. If the driver weaves too much, the score decreases. If the score drops below a threshold, the driver is giving a message to get rest. At a lower level, AutoTrak also provides a constant graphical display of drivers' lateral position. There is anecdotal evidence that having this type of display improves the ability of people to drive vehicles substantially different than what they are used to; a sedan driver who is driving a mini-van, for example. Improving driver feedback methods could result in improved driver performance, which would lead to fewer alarms (nuisance or otherwise), and greater system effectiveness and acceptance.

5.3 Summary

This thesis presents a novel warning algorithm, alarm decision model, and a methodology for training parameters of this model for an individual driver. The goal of this thesis, which is to reduce nuisance alarms while maintaining adequate warning time, through the use of an improved warning algorithm and training methodology has been met for certain types of drivers. There is still an open question regarding the statistical significance of the results, due to lack of data; so I provide the outline of a future research plan to determine the significance of the results. The data, experimental methodology, and analysis presented here, will, I hope, lead to more effective lane departure warning systems, and reduce the number of run-off-road crashes.

Appendix A Rural Roads

A.1 Introduction

In this appendix, I present an analysis of warning system performance on the entrance of a rural section of curved road, and demonstrate that, providing adequate lane departure warning in such situations is very difficult. While 40% percent of run off road incidents occur on rural roads [60], driver reaction time and the sharp curves present on rural roads combine to preclude developing a system which both has a low nuisance alarm rate and warns early enough to prevent a roadway departure.

There are many reasons for this. Rural roads are often narrow, have only one lane in each direction, curve sharply, and have either narrow or unpaved shoulders. Highways, in contrast, can be multi-lane, generally have 3.6m wide lanes, curve relatively gently, and often have four to five meter shoulders. The presence of such a shoulder allows the driver to briefly depart the road with no negative consequences. This, in turn, allows the warning system to give the driver more leeway, letting the vehicle approach closer to the lane boundary before sounding a warning. A rural road with no paved shoulder, in contrast, gives the driver very little latitude for deviation, especially on the entrance to a curve. A warning system in this environment would have to trigger when the driver is quite far from the lane boundary. This increased sensitivity would result in numerous nuisance alarms, possibly prompting the driver to ignore the system, or turn it off.

This analysis begins with the most optimistic set of parameters, which demonstrates that even when all conditions are ideal, preventing lane departures on rural roads near curves is challenging. The set of parameters includes:

1. A perfect driver, who is always centered in the lane, follows the road curvature and heading exactly, and drives at the posted speed limit.
2. A road in which curves have entry spirals which smoothly change as a clothoid.
3. A perfect, clairvoyant warning system which triggers the moment the driver deviates from the center of the road.
4. A driver who reacts to a warning in one second.

Using these optimistic parameters, I show that if an alarm occurs near the entrance to a curve, the driver will be more than a third of the way to the lane boundary before having a chance to react. Following this "ideal" analysis is a section which describes the same analysis conducted using roads which are not quite clothoidal, and where driver behavior is more realistic. This analysis will show that under less than ideal conditions, preventing lane departures on rural roads is intractable.

Following the analysis of driver performance on curves, I present a calculation of expected alarm rates on rural roads, not limited to curves. This work makes use of approximately 1 hour of real world data collected on a stretch of rural road near Columbus, Ohio, separated over 17 drivers, and shows that the presence of an adequate shoulder is vital to reducing nuisance alarm rates to an acceptable level.

A.2 Curve Analysis

This section presents the first "ideal" analysis mentioned in the introduction. The section starts with a discussion of the ideal and realistic driver and road models used, and presents results in predicting lateral deviation on the entrance to a curve. In this analysis, I demonstrate that if all conditions are ideal, it is possible to warn a driver in time to prevent a lane departure. However, Section A.2.3 shows that if these ideal conditions are relaxed even slightly, then the limitations of warning system performance and driver reaction time make it extremely difficult to warn the driver in time to prevent a lane departure while, at the same time, maintaining a low nuisance alarm rate.

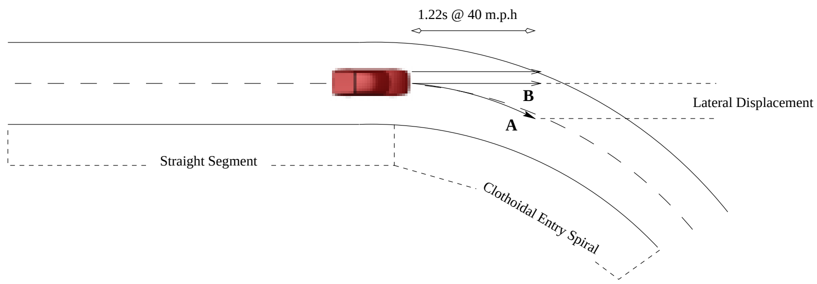

Figure A-1: A Perfect Driver -- vehicle heading, curvature, and position match the road perfectly. "A" is the path a perfect driver would follow. "B" is the path an inattentive driver would follow.

A.2.1 Ideal Analysis

This first part of the analysis shows that, given ideal conditions, warning a driver in time to prevent a road departure is feasible. This section presents descriptions of the perfect driver and road models which I use, followed by the methodology and results.

A.2.1.1 Perfect Driver Model

As Section 1.3 shows, there are many different ways to model a driver. For the first part of the analysis, a simple driver model is most appropriate. This model is that of a "perfect" driver. A perfect driver is one who follows the instantaneous heading and curvature of the road exactly, while staying perfectly centered within the lane. A perfect driver also travels at the posted speed limit. Figure A-1 demonstrates this. When the driver is tracking the road, the vehicle's path is locally a circular arc, where the curvature of the arc is equal to the instantaneous curvature of the road. This perfect path is marked as path "A".

When the warning system triggers an alarm, the driver has a one second reaction time before being able to react. This reaction time is simulated by freezing the vehicle curvature (for the duration of the reaction time) to that of the road curvature where the alarm is triggered. At this point, the driver is no longer perfectly tracking the road, and instead follows a circular arc for the duration of the reaction time. This path is shown as path "B". The distance from the endpoint of path “B” to the road centerline is the lateral deviation of the driver. The equations describing the (x, y) position of the vehicle in this state, after traveling a given distance *s* are:

$$x(s) = x_0 + \frac{\sin(cs + \theta_0)}{c} \tag{A-1}$$

$$y(s) = y_0 - \frac{\cos(cs + \theta_0)}{c} \tag{A-2}$$

Where *c* is the curvature of the arc, θ_0 is the initial heading, (*x0*, *y0*) is the (x, y) location of the center of the circle, and (*x(s)*, *y(s)*) is the (x, y) location of point *s*.

A.2.1.2 Perfect Road Model

A test subject collected data for this analysis by driving over a section of rural road while AutoTrak recorded longitudinal velocity, *v*, and yaw rate, y_r. The subject driving the vehicle took care to steer along the road as closely as possible, to avoid over- and under-steering, and to stay near the posted speed limit of 35 m.p.h. The instantaneous curvature of the road, *C*, can be computed using the relation. $C = (2\pi y_r)/(360v)$. Figure A-2 shows the curvature of the road R_{test}, which was used for this analysis. The top graph shows the curvature over the entire road, where positive curvatures are rightward turning roads, and negative curvatures are leftward turning. The bottom graph shows a zoomed in version of the area between the horizontal lines. This curve was chosen for analysis based on its sharpness (the peak radius of curvature is 66 meters). The portion of interest is in the range (2050, 2100) meters, which will be denoted C_R. This portion of the curve is known as the entry spiral.

According to highway design specifications, the entry spiral to a curve should be clothoidal, with the curvature increasing linearly over path length. This spiral should then be followed by a constant curvature circular arc, leading to an exit spiral which is also clothoidal. This configuration allows for a smooth traversal of the curve, as the driver has to turn his or her steering wheel at a constant rate to enter and leave the curve.

For the ideal analysis, I use this road model. The equations describing the clothoidal entry spiral are:

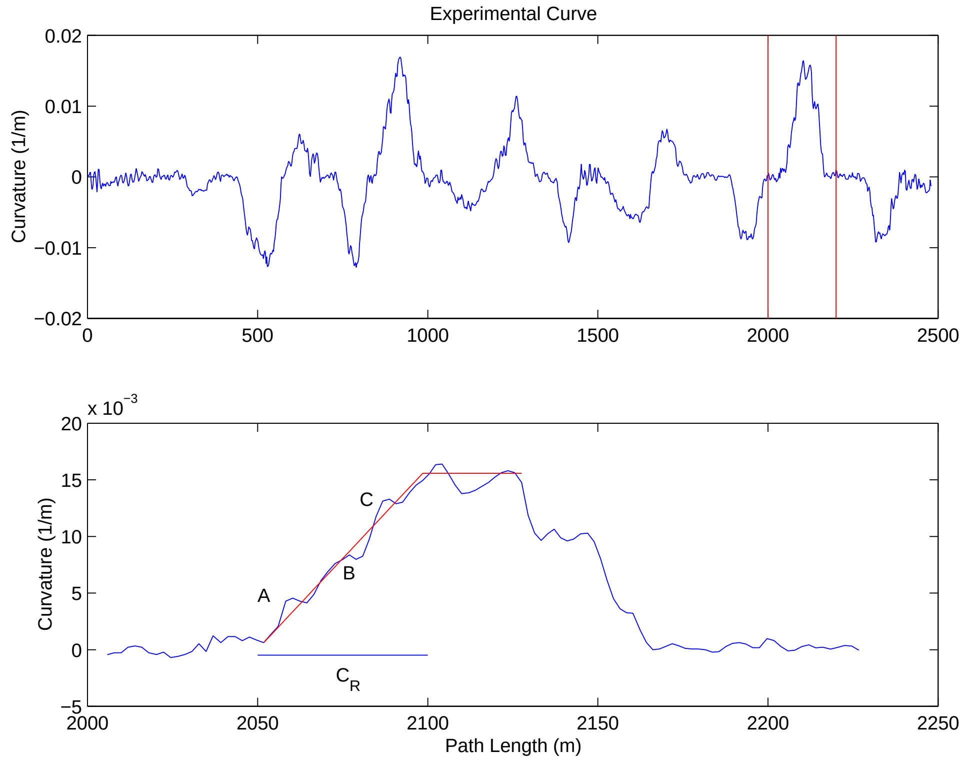

Figure A-2: Graph of the curvature used for analysis. The top graph is the curvature over the entire road; the bottom graph is a close-up of the region between the two horizontal lines. The trapezoid in the bottom graph is a least squares clothoidal approximation of the curve. Points A, B, C are areas where the radius of curvature error is greater than 20 meters.

$$c(s) = ks + c_o \qquad \text{(A-3)}$$

$$x(s) = x_0 + \int_0^s \cos\left(\frac{kt^2}{s} + c_0 t + \theta_0\right) dt \qquad \text{(A-4)}$$

$$y(s) = y_0 + \int_0^s \sin\left(\frac{kt^2}{s} + c_0 t + \theta_0\right) dt \qquad \text{(A-5)}$$

Where k is known as the sharpness of the clothoid, c_0 is the initial curvature, θ_0 is the initial heading, and $c(s)$, $x(s)$, and $y(s)$ are the curvature, x-position, and y-position of the clothoid, respectively, parameterized over the path length, s [33]. As equations A-4 and A-5 cannot be solved analytically, they are evaluated numerically using Fresnel integrals [27].

I compute the sharpness of the clothoidal entry region C_R using a linear least squares fit to curvature data in C_R. The slope of this line is k. This linear approximation to the curvature data is shown in red in Figure A-2.

A.2.1.3 Results

The goal of this analysis is to compute the expected lateral deviation which would result if given the following assumptions: 1) a perfect driver, 2) a clothoidal section of road, 3) a one second reaction time to a lane departure warning, and 4) a warning system which triggers the moment the driver is no longer perfectly tracking the road. The previous two sections described the perfect driver and road models which I use for this analysis. This section describes the methodology and results of the analysis.

The road, R, is a 50 meter long clothoidal path, followed by a 35 meter long circular arc, where the sharpness of the clothoid is based on the linear fit to region C_{R} in A-2 on page 113. The vehicle, V, starts at point s meters along the path, where $s \in (0, 50)$. Up to point s, the vehicle is perfectly tracking the position, heading, and curvature of the road, similar to path "A" in Figure A-2. The driver is traveling at 35 m.p.h., which is the posted speed limit of the curve. The warning system triggers an alarm at s. At this point, I simulate the effect of the driver's one second reaction time by freezing the curvature of his path, so that the vehicle is no longer perfectly tracking the road. This new circular path, which is similar to "B," deviates from the road. Before the driver has a chance to react and correct, the vehicle will have travelled a distance proportional to his velocity, and will have deviated from the center of the road. Equations A-2, A-3, A-4, and A-5 can be used to compute this deviation, as:

$$D = min\left(\sqrt{(R_x - V_{x(s)})^2 + (R_y - V_{y(s)})^2}\right) \qquad \text{(A-6)}$$

Where D is the deviation, $(V_{x(s)}, V_{y(s)})$ is the final location of the vehicle, and (R_x, R_y) is the clothoidal road. Figure A-3 shows a plot of this distance, where the x-axis indicates that point along the path where the alarm occurred, and the y-axis is the calculated deviation. The figure shows that a perfect driver on a perfect road with a perfect warning system will drift 0.21 meters before having a chance to react. Assuming a 3.0 meter road, and a 1.8 meter vehi-

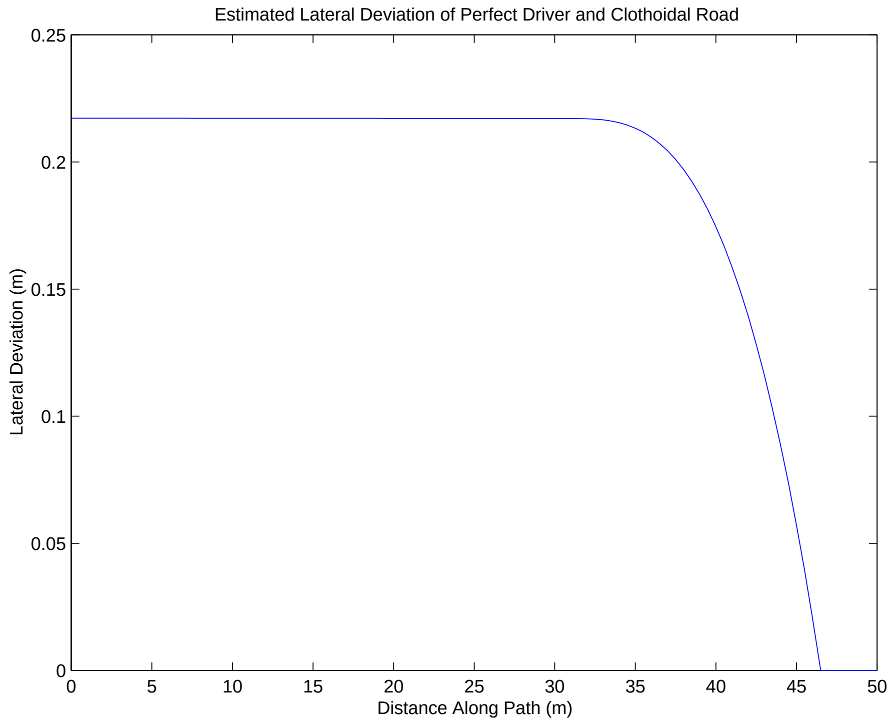

Figure A-3: Estimated lateral deviation for a perfect driver on an ideal road.

cle, the driver has 0.6 meters of leeway on each side. Therefore, a 0.21 meter drift is small enough that the driver would be able to recover. Note that the decrease in deviation which occurs when the warning system triggers beyond 35 meters along the curve is due to the presence of the circular arc, where the curvature is constant. If the perfect driver gets a warning while on the circular arc, he will continue to follow the arc, and track the road.

A 0.21 meter drift gives hope that preventing accidents on rural curves is possible. However, this result assumes a clairvoyant warning system. A realistic warning system would react when the driver is closer to the lane boundary, cutting into the allowable deviation.

A.2.2 Ideal Time to Collision

This section uses the models developed above to present a short calculation which shows that, at the entrance of a curve, an inattentive driver travelling at 40 m.p.h would cross a lane boundary in 1.22 seconds. If the lane width is 3.0 meters, and the vehicle width is 1.8 meters, a lateral displacement of 0.6 meters would leave the driver with one wheel touching a lane boundary.

Evaluating Equation A-6 until $D=0.6$ shows that the left tire of a vehicle travelling along path “B” at 40 m.p.h. (17.77 meters/sec.) would touch the lane boundary in 1.22 seconds. This means that, if there is no shoulder, the entire warning process, from detection of impending departure to the completion of driver correction, must occur within 1.22 seconds.

A.2.3 Realistic Analysis

The first analysis demonstrates that, given the ideal conditions outlined in Section Section A.1, it *may* be possible to warn drivers of impending lane departures while they still have enough time to prevent them. This analysis however, is unrealistic. The entrances to curves are generally not perfectly clothoidal, and real drivers definitely do not behave as the previous “perfect” model suggested. Instead, they make errors when tracking a road. In this section, I will show that these errors, combined with a more realistic road model, make early detection of lane departures in curves intractable. I still use a clairvoyant warning system, which triggers the moment the driver stops perfectly tracking the road.

A.2.3.1 Realistic Road Model

In theory, roads are designed with a clothoidal entry spiral followed by a circular arc. In practice, issues such as cost, ease of construction, and the local landscape affect the design of a road. These issues often preclude using a clothoidal spiral in a curve, particularly in rural areas. Therefore, it is difficult to know how to properly model a curve without having access to “as-built” blueprints, which show what was *really* built, as opposed to what was designed.

On inspection, the region C_{R} in Figure A-2 appears to be clothoidal. The bottom figure shows the results of a least squares clothoidal fit to the entry and constant curvature portion of the curve. Empirically, the fit looks good. However, closer inspection reveals that there are ranges along C_{R} where the radius of curvature error exceeds 20m. These points are marked “A,” “B,” and “C” on the figure. This could be due to over- or under-steering, but it could also be due to actual changes in the curvature of the road. The road represented by R_{test} was driven over three times. In each case, the region C_{R} showed similar variations from linearity. This implies that the non-linearity is a feature of the road, and not necessarily an artifact introduced by the driver.

Therefore, I use a local approach to model the road curvature at a given point. The road curvature R$C_{(s)}$, at point *s*, (where $s \in C_R$) is the average of the radii of curvature at $C_{(s-(d/2))}$ and $C_{(s+(d/2))}$. *d* is the path length a driver would traverse during a reaction time of *t* seconds at a velocity *v*. For this analysis, the reaction time *t* is always 1.0 second. Intuitively, this model assumes the road will follow a circular arc whose radius of curvature is the average of the radii of curvature at the start and end points of the path to be approximated. The length of the path to be approximated is the distance the driver would traverse during one second of reaction time. Therefore:

$$RC_{(s)} = \left(\frac{C_{\left(s-\left(\frac{d}{2}\right)\right)}^{-1} + C_{\left(s+\left(\frac{d}{2}\right)\right)}^{-1}}{2}\right)^{-1} \qquad \text{(A-7)}$$

A.2.3.2 Realistic Driver Model

The perfect driver model in Section A.2.1.1 does not exist. Drivers do not track the position, heading, and curvature of a road perfectly. This “realistic” driver model I use in this part of the analysis relaxes these assumptions. This driver still tracks the center of the road. However, he or she is allowed to oscillate a bit. This oscillation is manifested as a 1 degree error in the yaw of the vehicle relative to the road, simulating the driver “driving through” the center of the road at a slight angle while oscillating. The driver also travels at 5 m.p.h. over the posted speed limit, which is 35 m.p.h.

A.2.3.3 Results

The previous sections presented a more realistic driver and road model. The two models simplify down to circular arcs of differing curvatures. This makes the analysis of the potential displacement very convenient. This section presents the results on what the expected deviation would be for a realistic driver with a one second reaction time on a locally circular road. This deviation is computed by looking at the distance between the end points of two circular arcs; one representing the driver, and the other representing the road.

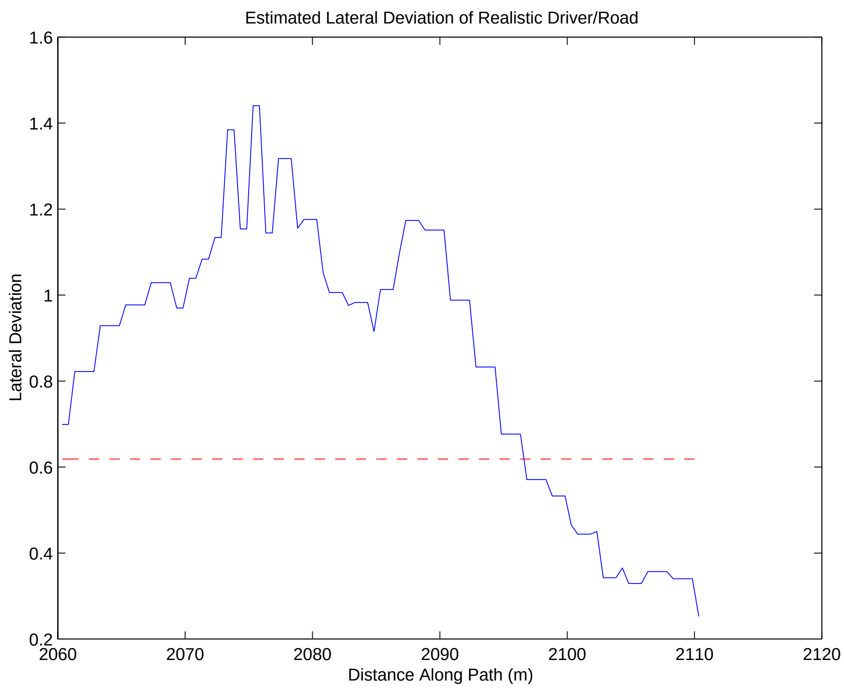

Figure A-4: Expected lateral deviation of a realistic driver on a realistic road. The dashed line is the expected deviation of a realistic driver on a perfect road.

Figure A-4 shows the results of computing the expected deviation over the first 30 meters of C_R. The x-axis is the location in the entry spiral where the alarm goes off and the driver freezes. The y-axis is the expected deviation, which is computed using Equation A-7 for the road model. The solid line is the expected deviation computed using a realistic road and driver model.

The dashed line is the expected deviation if the entry spiral C_R is modeled as a clothoid, and is provided for contrast. We can see that the expected deviation is constant over the distance along the curve. This does not fit with what is observed in the realistic road model, further indicating that the road which was used for the test is not clothoidal. The deviation is greater than 0.6m, indicating that the vehicle would depart the lane. This means that, even with a perfect road, a driver who is traveling 5 m.p.h. over the speed limit and has a one degree error in heading will exceed the lane boundary.

The realistic results show that the driver can end up with a tire at least 1.4 meters over the lane boundary. This higher deviation when compared with the clothoidal road is due to the occasional sharp changes in the curvature of the road, combined with the vehicle heading error. A 1.4 meter deviation means that the outside tire is 0.8m beyond the boundary before the driver has a chance to react. If the road has a sufficiently wide shoulder, then the driver has a chance to correct. However, as many rural roads do *not* have wide shoulders, the vehicle would end up departing the road. Again, keep in mind that these results are based on a clairvoyant warning system. A realistic warning system would result in even larger deviations.

A.2.4 Summary

Given a perfect driver (who is traveling 5 m.p.h. above the speed limit), a clothoidal rural road, and a clairvoyant warning system, the following events have to happen in 1.22 seconds: 1) The warning system has to detect that the driver is drifting. 2) The warning system has to sound an alarm. 3) The driver has to perceive the warning. and 4) The driver has to take corrective action to avoid the road departure.

This result, however, assumes a *perfect* driver. A normal driver is almost never perfectly centered in the lane. Similarly, drivers do not perfectly track the curvature and heading of a road while entering a curve. More importantly, lane departure detection systems are not clairvoyant. They do not trigger alarms when the driver is very near the center of the lane, accurately tracking the road. Rather, they trigger when the driver is off center, and heading off road. This means that the warning system really has *less* than one second in which to complete the four steps outlined above, making the problem not only unreasonable, but intractable.

A.3 Rural Road Alarm Rates

In the previous section, I present the expected lateral deviation given various driver and road models, and a clairvoyant warning system. It concentrates on the performance of a warning system on the entry to a curve, which can be a particularly treacherous section of road to navigate in rural areas. This section examines the performance of a realistic warning algorithm on real data collected from 17 drivers who drove over a stretch of rural road. This stretch of road includes both curves and straight segments.

The section begins with a brief description of the data, followed by the methodology of the experiment. The results show that, for rural roads with less than a meter of shoulder, the number of alarms triggered per hour becomes unacceptably high.

A.3.1 VRTC Data

The data used in this experiment was collected by the members of the Vehicle Research Test Center (VRTC) in Columbus, Ohio, using their DASCAR data collection vehicle [14]. This vehicle uses downward looking cameras mounted on the sides of the rear of the vehicle to track lane boundaries and derive lateral position. Sixty-four subjects drove the DASCAR on a series of rural and interstate roads while data on their lane keeping performance was logged. Of the 64 subjects, 17 produced enough valid data on the rural section of road to be of use in this analysis.

A.3.2 Methodology

I applied a first order future offset distance (FOD) warning algorithm to each of the 17 rural road datasets. The FOD algorithm, which is discussed in Section 3.3, predicts the lateral position of a vehicle FOD_t seconds into the future, where FOD_t is the lookahead time. If the predicted lateral position is more than FOD_d meters beyond the lane boundary, a warning is triggered. FOD_d is the virtual lane boundary. Setting FOD_d greater than 0 meters has the effect of allowing the driver to wander beyond the lane boundary by FOD_d meters. For this experiment, the prediction was generated using a first order model, $L_p' = L_p + (L_v \times FOD_t)$ where L_p' is the predicted lateral position, L_p is the current lateral position, and L_v is the lateral velocity.

When applying the FOD algorithm to the 17 datasets, the lookahead time parameter was varied over the range [0, 2.0] seconds in increments of 0.2 seconds. The offset distance parameter was in the range [0.0, 0.6, 1.2, 1.8] meters (corresponding to [0, 2, 4, 6] ft.). Each pair (FOD_t, FOD_d) was applied to each dataset. The alarms per hour for each run were stored, and averaged over drivers. This resulted in one alarm/hour metric for each (FOD_t, FOD_d) pair.

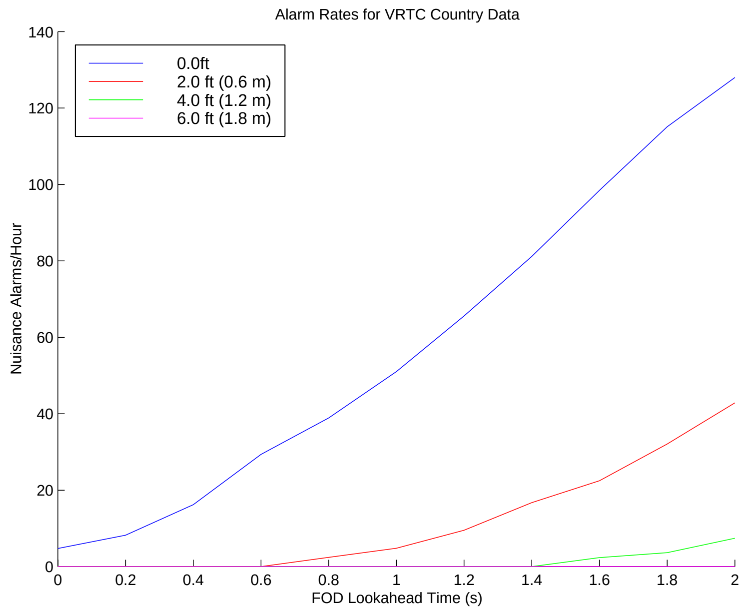

Figure A-5: FOD Alarm Rates of VRTC Rural Data

Figure A-5 shows the results of this experiment. The x-axis is FOD_t and the y-axis is the number of nuisance alarms per hour. Each graph is a different FOD_d. Concentrating on (FOD_t = 1.0s), the results show that (FOD_d = 0.0m) results in the highest alarms/hour. This is not surprising, as (FOD_d = 0.0m) is equivalent to triggering an alarm anytime the predicted lane position is beyond the lane boundary, which is appropriate when there is no paved shoulder. Even with (FOD_d = 0.6m), the alarm rate at (FOD_t = 1.0) is greater than 5 nuisance alarms/hour. An alarm rate this high would either annoy drivers and prompt them to disable the system, or habituate them to the alarm sound, reducing effectiveness. It is only when the allowed deviation is greater than 1.2m does the alarm rate drop down to 0. As this dataset had no lane changes which could look like valid road departures, and no accidents, the alarm rate should be zero.

Note that nuisance the alarm rate at (FOD_t = 0.0s, FOD_d = 0.0m) is approximately 6 alarms/hour. This means that the driver exceeded the lane boundary at this frequency, during normal driving. However, the road which the subjects were driving had a 4-5ft shoulder, which means that the deviations did not place the vehicle in jeopardy. Presumably, drivers on a rural road with no shoulder would drive “tighter”.

A.4 Conclusion

This appendix presents an analysis demonstrating that even given a perfect driver on a simple road with a clairvoyant warning system, accurately predicting lane departures at the entrance to a curve while maintaining an acceptable nuisance alarm rate is very difficult. This analysis is then expanded to include realistic road and driver models. The large lateral deviations indicate that effectively preventing road departures on rural curves is exceedingly difficult. This was further demonstrated using real data collected by the VRTC, which showed that alarm rates become unacceptably high when trying to predict the lateral position of a driver far enough into the future so as to give him or her enough time to appropriately react to a warning.

This does not, however, mean that the problem is *impossible* to solve for all rural roads. There is evidence that 40 percent of rural road accidents occur on roads with shoulders which are narrower than 3 feet. If an adequate shoulder is present, Figure A-5 demonstrates that the nuisance alarm rate is zero. However, a passive warning system would have to know the shoulder width of the road on which it was operating, to decide whether it is capable of detecting lane departures in time to prevent them. While such work is beyond the scope of this thesis, annotated maps are one possible way of obtaining such information. Similarly, this entire appendix has dealt only with passive warning systems. Active warning systems have either steering or differential braking capability to assist the driver in staying close to the center of the lane. Such systems could conceivably reduce the rate of rural road run off road incidents even on roads with narrow or no shoulders, although this has not been demonstrated.

APPENDIX B **HURP Documents**

This appendix includes the proposals required by the Carnegie Mellon University Human Use Review Panel, and refer to the studies described in Section [REF]. They are presented as templates for use by researchers at other institutions who wish to conduct similar studies. The first section is on the NavLab 8 study, and includes the HURP proposal, consent form, subject questionairre, and post-experiment questionairre. The second section is on the naturalistic driver study, and includes the HURP proposal and consent form.

Proposal for Driver Characteristics Human Studies

Parag H. Batavia, Dean A. Pomerleau, Charles E. Thorpe

Abstract

This study proposes to collect driving performance data from 20 individuals, spread by gender and age. The data will be collected while they are on extended trip in their personal vehicle, which will be out-fitted with a data collection system.

The purpose of this study is to develop models of highway driving behavior that can be used to tailor unexpected lane departure warning systems to individual users, increasing accuracy and decreasing false alarms. A unexpected lane departure is defined as any human or automobile action or trajectory which causes the vehicle to unexpectedly leave the lane that it is currently in. A controlled lane change would not deserve a warning, whereas drifting due to fatigue or inattention would.

Current research in lane departure warning systems ignores the individual traits of the driver. Most methods use only physics-based models, which look at the current trajectory of the vehicle, and calculate a ‘time to lane crossing’, which is measured in seconds. When this number falls below a certain threshold, a warning is sounded. This type of system does not take into account the normal tendency of a given driver to weave slightly, or prefer one side of a lane to another. By ignoring these types of factors, the warning system has to be tuned to the lowest common denominator, which causes an unacceptable false alarm rate. Research has shown that if the warning system cries wolf too many times, the driver tends to ignore it. It is hoped that this work will not only reduce the number of false alarms, but increase the accuracy and improve the reaction time of the warning.

However, to build these types of warning systems, a large amount of data is needed to determine what types of different traits drivers have, and how they can be exploited. For this reason, we wish to collect driving data on a large number of individuals. The data will be collected by the RALPH lane tracking system, which will passively monitor and record the

driving behavior of the subjects. The data collected will include current position in the lane, the driver's steering wheel position, velocity, and information on surrounding vehicles measured with on-board radar and laser. A video tape recording of the road ahead will also be made via a camera mounted under the rear-view mirror, pointing outwards.

Subject Utilization

The 36 individuals selected for this experiment will vary in age from 21 (the minimum age required by CMU insurance to drive a CMU owned vehicle) to 55, and will be evenly split by gender. The subject will be asked to fill out a questionnaire, which is attached. Subjects will be asked to sign a liability waiver and submit a copy of their driver's license for verification of validity, as per CMU insurance policy. The subject will be asked to drive Navlab 8, an Oldsmobile Silhouette van, from CMU to a destination approximately 1 hour away, and then back. He will be told that the purpose of the trip is to collect data, but besides that, he will not be told what is precisely being monitored.

At all times during the trip, the experimenter will be present in the passenger side. The experimenter will not interfere with the subject except in two cases: 1) if the subject is driving at a dangerous speed, or otherwise erratically, and 2) at random times during the trip, the experimenter will ask the subject to perform a lane change. The subject will be asked to perform this lane change *whenever he deems it is safe to do so*. If the driver feels uncomfortable making the lane change, he will have complete freedom to disregard the instruction. If at any point during the trip, the driver feels that he is unable to continue, the experimenter will return the subject to CMU. Furthermore, the trip will only be made if road conditions are clear throughout the intended route, and no precipitation is forecast.

Sample Consent Form

See Attached

Confidentiality

Confidentiality will be maintained by assignment of a number to each driver. The only driver-specific information that will be recorded is what is on the questionnaire. All computer files and video tapes will be labeled only with the code number.

Risk/Benefit Analysis

There is no direct benefit to the subjects participating in this experiment. There is, however, a benefit to society. Currently, approximately 90% of all traffic accidents are caused by driver error. This includes nearly 14,000 deaths each year due directly to run-off-road incidents, which is 1/3rd of all driving fatalities.A lane departure warning system will aid in alleviating this form of error. By reducing the number of false alarms, the system will be more easily accepted by the driver. Increasing the accuracy and using driver specific information will decrease the warning time and increase the reaction time available to the driver, preventing possibly fatal lane departures. The risk to the participant is negligible, and equivalent to the normal risk assumed when operating a vehicle on a 2 hour trip.

Carnegie Mellon University

Driver Characteristics Study
Conducted By The Robotics Institute
Consent Form

I agree to participate in experimental research conducted by members of the faculty or by students under the supervision of members of the faculty. I understand that the proposed research has been reviewed by the University's Institutional Review Board and that to the best of their ability they have determined that the observations involve no invasion of my rights of privacy, nor do they incorporate any procedure or requirements which may be found morally or ethically objectionable. If, however, at any time I wish to terminate my participation in this study I have the right to do so without penalty. I have the right to request and keep a copy of this form.

If you have any questions about this study, you should feel free to ask them now or anytime throughout the study by contacting:

Dr. Charles E. Thorpe
Robotics
224 Smith Hall
(412)268-3612
cet@ri.cmu.edu

You may report any objections to the study, either orally or in writing to:

Susan Burkett
Associate Provost
Carnegie Mellon University
(412)268-8746

Purpose of the Study: I understand that I will be driving an Oldsmobile Silhouette Minivan on a round trip from CMU to a destination one hour away. I realize that data about my driving performance will be recorded, and a video will be kept of the road in front of me. I will not be able to be identified by this video. If at any time I do not wish to continue with the experiment, the experimenter will return me to CMU with no questions asked and no penalty.

I understand that the following procedure will be used to maintain my anonymity in analysis and publications/presentations of results. Each participant will be assigned a number. The researchers will save the data and videotape files by participant number, not name. Only members of the research group will view the tapes in detail.

I understand that in signing this consent form, I give Professor Thorpe, and his associates, permission to present this work in written and oral form without further permission from me.

______________________	______________________
Signature	**Date**
______________________	______________________
Print Name	**Telephone**

Consent for public display of experimental videotape

I give my permission for the video tape of the experiment in which I have participated to be shown in public. I realize that if the tape is shown in public, viewers will not be able to identify me. I have the right to refuse this request without penalty.

______________________	______________________
Signature - I give permission	**Signature - I refuse permission**

Number: ______

Driver Characteristics Study Questionnaire

General

Age:____________________ Gender: __ Male __ Female

Driving Experience

Do you have a driver's license? ___________________

How long have you been licensed? ___________________

If yes, how often do you currently drive? ___________________

What kind of car do you normally drive? ___________________

Have you had a moving violation within the last year? _______

Number: ______

Driver Characteristics Post-Study Questionnaire

For each question below, circle the number that best indicates how you feel. The numbers range from 1 to 5, with 1 meaning 'I don't agree at all', and 5 meaning 'I agree completely'

	Disagree				Agree
I felt uncomfortable driving a mini-van	1	2	3	4	5
Having someone else in the car made me nervous	1	2	3	4	5
My driving style was different today than normal	1	2	3	4	5
I tried to stay in the center of the lane	1	2	3	4	5
I had trouble remembering to use the turn signal.	1	2	3	4	5
I had trouble judging where the right side of the vehicle was	1	2	3	4	5
I drove slower than I normally do	1	2	3	4	5

Proposal for Driver Characteristics Human Studies

Parag H. Batavia, Dean A. Pomerleau, Charles E. Thorpe

Abstract

This study proposes to collect driving performance data from 20 individuals, spread by gender and age, on an extended trip (> 6 hours) or daily commute in which the subject's vehicle is out-fitted with a data collection and lane departure warning system.

The purpose of this study is to analyze the performance of our lane departure warning system. This system uses a video camera mounted under the rear-view mirror and an associated computer system to scan the road ahead of the vehicle, and determine the lane position of the car. This lane position is then used to extrapolate the future position of the vehicle. If it appears as if the driver is about to depart the road (either due to inattention, drowsiness, or other impairment), the system sounds an alarm to alert the driver to pay attention (although, as we will describe below, this audible warning will not be used in this study).

We wish to study the performance of the system in two areas: system availability, and warning accuracy. System availability will involve logging information about when the image from the video camera is sufficient for our computer to use to detect lane markings. Situations such as heavy fog, high glare, hard rain, and lack of road markings can cause the system to not function properly. We wish to determine how often these situations occur.

Furthermore, we wish to analyze the accuracy of our warning algorithm. This is the part of the system which looks at the vehicle's lane position, and determines if the driver is in danger of departing the road. To this end, we will be logging data on the state of the vehicle, such as its lane position, velocity, and the curvature of the upcoming road. We will also log events when the warning system would normally trigger a warning. However, in this initial study, no feedback will be provided to the driver.

From the subject's perspective, the system will be totally silent and nearly unobtrusive. An interface unit which contains a camera will be mounted (using suction cups) under the rear-view mirror, and this unit will be connected to a computer system and VCR. The computer system and VCR will either be placed under the passenger seat, or in the rear of the vehicle.

Subject Utilization

The 20 individuals selected for this experiment will vary in age from 18 to 65, and will be evenly split by gender. The subject will be asked to fill out a questionnaire, which is attached. Subjects will be asked to sign a liability waiver, which is also attached. The subject will be asked to bring his vehicle to CMU so that we can install the system. We will then provide a briefing to the subject, telling them how to activate and de-activate the system. After the subject returns from her trip, she will be asked to return to CMU so we can remove the lane tracking system. The subject will then be provided with a page describing the purpose of the study, which is attached.

Sample Consent Form

See Attached

Confidentiality

Confidentiality will be maintained by assignment of a number to each driver. The only driver-specific information that will be recorded is what is on the questionnaire. All computer files and video tapes will be labeled only with the code number.

Risk/Benefit Analysis

There is no direct benefit to the subjects participating in this experiment. There is, however, a benefit to society. Currently, approximately 90% of all traffic accidents are caused by driver error. This includes nearly 14,000 deaths each year due directly to run-off-road incidents, which is 1/3rd of all driving fatalities. A lane departure warning system will aid in alleviating this form of error. By reducing the number of false alarms, the system will be more easily accepted by the driver. Increasing the accuracy and using driver specific information will decrease the warning time and increase the reaction time available to the driver, preventing possibly fatal lane departures. There is no increase in risk to the subjects, since they will be driving on trips they would be taking anyway, and the system will simply be unobtrusively recording their driving behavior.

Carnegie Mellon University

Driver Characteristics Study
Conducted By The Robotics Institute
Consent Form

I agree to participate in experimental research conducted by members of the faculty or by students under the supervision of members of the faculty. I understand that the proposed research has been reviewed by the University's Institutional Review Board and that to the best of their ability they have determined that the observations involve no invasion of my rights of privacy, nor do they incorporate any procedure or requirements which may be found morally or ethically objectionable. If, however, at any time I wish to terminate my participation in this study I have the right to do so without penalty. I have the right to request and keep a copy of this form.

If you have any questions about this study, you should feel free to ask them now or anytime throughout the study by contacting:

Dr. Dean A. Pomerleau
Robotics Institute
211 Smith Hall
(412)268-3210
pomerlea@ri.cmu.edu

You may report any objections to the study, either orally or in writing to:

Susan Burkett
Associate Provost
Carnegie Mellon University
(412)268-8746

Purpose of the Study: I understand my vehicle will have an AutoTrak driver performance monitor system installed. I realize that data about my driving performance will be recorded, and a video will be kept of the road in front of me. I will not be able to be identified by this video. If at any time I do not wish to continue with the experiment, I may turn off the recording device and return to CMU.

I understand that the following procedure will be used to maintain my anonymity in analysis and publications/presentations of results. Each participant will be assigned a number. The researchers will save the data and videotape files by participant number, not name. Only members of the research group will view the tapes in detail.

I understand that in signing this consent form, I give Professor Pomerleau, and his associates, permission to present this work in written and oral form without further permission from me.

Signature **Date**

Print Name **Telephone**

Consent for public display of experimental videotape

I give my permission for the video tape of the experiment in which I have participated to be shown in public. I realize that if the tape is shown in public, viewers will not be able to identify me. I have the right to refuse this request without penalty.

Signature - I give permission **Signature - I refuse permission**

APPENDIX C NavLab 8 Local Disturbance Evidence

C.1 Introduction

In Section 3.6.2, I present an extension to the Future Offset Distance (FOD) based alarm decision model which takes local deviations in driver behavior into account. My justification for doing this is that changes in traffic, road environment, and driver behavior can all cause biases in lane keeping performance. For instance, it is not uncommon for drivers in the left lane who are passing a truck or other vehicle to drift to the left, to allow more space on the right side. This drift can cause nuisance alarms for a lane departure warning system. In this appendix, I present evidence for this drift, using the NavLab 8 data, described in Section 2.3. I show that many drivers, when passing vehicles, tend to drift away from the vehicle they are passing. There can be many other reasons for such local deviations, such as the presence of a construction zone, or changes in road structure or markings. However, I do not have data on these effects, and cannot include them in this analysis.

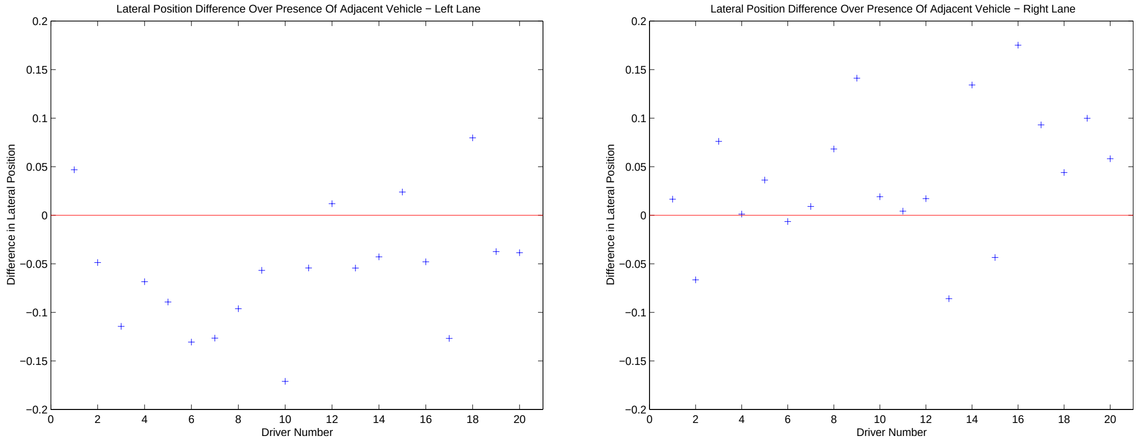

Figure C-1: Difference in mean lateral position with and without the presence of an adjacent vehicle for the 20 NavLab 8 subjects. The left graph is when the driver is in the left lane, the right graph is when the driver is in the right lane. Negative lateral positions and left of lane center, positive are right.

C.2 Evidence

The NavLab 8 data includes sensor information on the presence of vehicles in adjacent lanes. There are two radar-based rear side sensors, which detect the presence of vehicle (or other obstacles) on either side of the vehicle. This information can be used to calculate average lateral positions when adjacent vehicles are and are not present. I do this for two different cases: When the driver is in the left lane, and when the driver is in the right lane. When the driver is in the left lane, I calculate the average lateral position for when there is no adjacent vehicle, and for when one is present. If adjacent vehicles cause drivers to change their behavior, then the average lateral position in the presence of such a vehicle would be further left than if there was no adjacent vehicle, and the opposite is true when the driver is in the right lane.

Figure C-1 shows the difference in mean lateral position with and without the presence of an adjacent vehicle, for all 20 NavLab 8 subjects. The left graph is when the subject is in the left lane, and the right graph shows results for when the driver is in the right lane. The left graph shows that, for the majority of the cases, the mean lateral position with an adjacent vehicle is to the left of the mean lateral position in the absence of an adjacent vehicle. For a quarter of the subjects, the difference is 10cm or greater. The same is true when the subject is in the right lane. The majority of the subjects shift to the right when there is a vehicle to their left, and again, a quarter of them shift by 10cm or more. However, the subjects who shifted the most when in the left lane are not the same as the subjects who shifted the most in the right lane.

There is also greater consistency over subjects in the shifting behavior when drivers are in the left lane, opposed to when they are in the right lane. This could be because when drivers are in the left lane, they are more likely to be passing vehicles on the right, travelling at a higher speed, and therefore they may be more conscious of their position relative to the vehicle being passed. It could also be an effect of driving NavLab 8, which is a minivan. The next appendix shows that a few drivers had some trouble judging where the right side of the minivan was. This could make them more cautious when passing. Similarly, when in the right lane, drivers have a better feel for where they are relative to obstacles to their left, and may not feel as closed in as when they are in the left. In the majority of the cases, though, there is shifting behavior in both left and right lanes.

C.3 Summary

I have shown that in the presence of other vehicles, drivers tend to change their behavior. This change in behavior manifests itself as a change in lateral position, to move away from the potential obstacle. While the behavior seems to be a bit more pronounced and consistent when the driver is in the left lane, it occurs when the driver is in the right lane as well. Accounting for this shift improves warning system performance, as I demonstrate in Section 3.6.2.4.

APPENDIX D Navlab 8 Questionnaire

D.1 Introduction

This appendix looks at the results of the survey which subjects in the NavLab 8 study (described in Section 2.3) completed. This survey was completed by subjects after they had driven Navlab 8 and returned to CMU. The purpose of the questionnaire was to gauge the subject's perception of his driving performance, to use as a measure of possible confounding effects in the data. The subjects were asked to agree or disagree to the following statements, in a range from 1 to 5, where "1" was "disagree" and "5" was "agree":

1. I felt uncomfortable driving a mini-van.
2. Having someone else in the car made me nervous.
3. My driving style was different today than normal.
4. I tried to stay in the center of the lane.
5. I had trouble remembering to use the turn signal.
6. I had trouble judging where the right side of the vehicle was.
7. I drove slower than I normally do.

Figure D-1 shows a histogram of the responses to each of the questions.

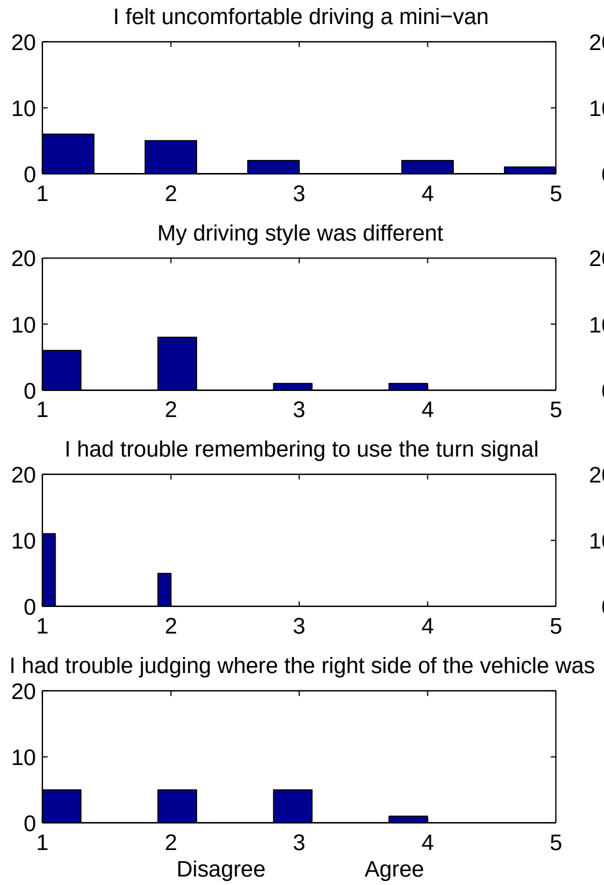

Figure D-1: Histograms of responses to the Navlab 8 post-study questionnaire.

D.2 Discussion

Section 2.3.3.2 discussed possible confounding effects which may have influenced the data collected in the Navlab 8 study. These effects included driving an unfamiliar vehicle, and the presence of an experimenter.

Question 1 addresses the level of comfort the subject experienced while driving a mini-van. Figure D-1 shows a histogram of responses to this question, and Table 4-1 shows the cross-correlation between the questions. Note that while the majority did not claim to feel uncomfortable, a significant percentage had responses of “4” or “5,” indicating that they felt uncomfortable driving the mini-van. However, more than half of the subjects claimed little or no discomfort in driving a mini-van.

Question	1	2	3	4	5	6	7
1	1.0	-0.25	-0.15	0.25	0.01	0.47	0.20
2	-0.25	1.0	-0.26	-0.45	0.38	-0.31	-0.24
3	-0.15	-0.26	1.0	0.46	0.16	0.28	0.13
4	0.25	-0.45	0.46	1.0	0.02	0.64	0.13
5	0.01	0.38	0.16	0.02	1.0	0.35	-0.16
6	0.47	-0.31	0.28	0.64	0.35	1.0	0.26
7	0.20	-0.24	0.13	0.13	-0.16	0.26	1.0

TABLE 4-1 Cross-Correlation Between Navlab 8 Survey Questions

My presence in Navlab as an experimenter was another possible confounding factor in subject behavior. However, the survey answers very strongly indicate that this is not so. The responses to question 2 are nearly all “1,” indicating that my presence did not affect driver behavior. This is not surprising, considering that most of the subjects were colleagues, whom I’ve known for a while. It is, instead, more likely that my presence actually had the opposite effect. Having an extra person in the vehicle, who is conversing with the driver, may help the driver forget that he is being monitored, and let him behave naturally. However, this is difficult to quantify, as no controlled experiments were done. Subjects also had no trouble remembering to use their turn signal, which was the only explicit instruction they were before the trial. This matches the observations in the data, which show that nearly all lane changes were marked with turn signals.

The responses to questions 1 and 2 suggest that while the differences between driving a mini-van and a sedan might have influenced driver behavior, my presence in the vehicle most likely did not. The remaining questions expand on these two possible sources of “contamination”.

Question 7, in particular, highlights one of the main differences between sedans and mini-vans, which is the difference in visibility out the right side. It is more difficult to judge the position of the right side of a mini-van, due to reduced visibility, than of a sedan. However, nearly all the respondents did not agree with question 7, indicating that they had little or no trouble judging the location of the right side of the mini-van. This is also indicated in the low correlation score (c = 0.20) between questions 1 and 7. This low correlation means that those drivers who did feel uncomfortable weren't necessarily uncomfortable due to restricted visibility. This restricted visibility did, however, have a possibly unconscious impact on their performance, as indicated by the leftward mean lane positions.

This disagreement could mean that another factor is to blame for the minor discomfort felt in the mini-van. There is a higher correlation between questions 1 and 6 (c = 0.47), which is whether the subject drove slower than normal. However, most subjects claimed that they did not slow down or temper their speed during the trial. Those that did were more likely to be those that felt uncomfortable in the mini-van. Question 3, which addresses the subjects' perceptions of their driving style, indicates that nearly all of the drivers felt they were behaving as they normally do.

D.3 Summary

This appendix has evaluated the perception of a Navlab 8 test subject vs. his actual performance, through the use of a post study questionnaire. The main conclusions which can be drawn from this analysis are that while vehicle type may be a confounding effect, the presence of an experimenter most likely was not. This entire analysis, however, was based on the subject's *perception* of his behavior. The small sample size (N = 17) makes the correlations shown in Table 4-1 difficult to judge. In any event, supposing causal relations due to correlation score is always risky.

References

[1] R. Wade Allen. Modeling driver steering control behavior. In *IEEE Proceedings of the International Conference on Cybernetics and Society*, pages 112–116. IEEE, 1982.

[2] R. Wade Allen. The driver's role in collision avoidance systems. Paper Prepared for the Workshop on Collision Avoidance Systems, 1994.

[3] P.E. An, M. Brown, and C.J. Harris. On real time driver modeling and vehicle guidance within prometheus. In *Proceedings of Transportation Systems: Theory and Technology of Advanced Technology*, pages 91–96, 1994.

[4] Christopher G. Atkeson, Andrew W. Moore, and Stefan Schaal. Locally weighted learning. *Artificial Intelligence Review*, 1997.

[5] Parag Batavia. Driver adaptive warning systems. Technical Report CMU-RI-TR-98-07, Carnegie Mellon University, March 1998.

[6] Parag H. Batavia, Dean A. Pomerleau, and Charles E. Thorpe. Predicting lane position for roadway departure prevention. In *Proceedings of the IEEE Intelligent Vehicles Symposium*, Stuttgart, Germany, October 1998.

[7] John Baxter and John Y. Harrison. A nonlinear model describing driver behavior on straight roads. *Human Factors*, 21(1):87–97, 1979.

[8] George A. Bekey, Gerald O. Burnham, and Jinbom Seo. Control theoretic models of human drivers in car following. *Human Factors*, 19(4):399–413, 1977.

[9] Erwin R. Boer, Marcon Fernandez, Alex Pentland, and Andrew Liu. Method for evaluating human and simulated drivers in real traffic situations. In *IEEE 46th Vehicular Technology Conference*, pages 1810–1815, 1996.

[10] 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.

[11] August Burgett. DOT's approach to its safety evaluation. In *Proceedings of the Safety Evaluation of Intelligent Transportation Systems Workshop*, pages 61–77, Reston, VA, May 1995.

[12] R.H. Byrne and C.T. Abdallah. Design of a model reference adaptive controller for vehicle road following. *Mathematical Computer Modelling*, 22(4-7):343–354, 1995.

[13] J.M. Carson and W.W. Wierwille. Development of a strategy model of the driver in lane keeping. *Vehicle Systems Dynamice*, 7:233–253, 1978.

[14] R.J. Carter, F.S. Barickman, and M.J Goodman. DASCAR sensor suite and video data system. In *Proceedings of the SPIE - Transportation Sensors and Controls: Collision Avoidance, Traffic Management, and ITS.*, volume 2902, pages 250–259, 1997.

[15] Robert E. Chandler, Robert Herman, and Elliot W. Montroll. Traffic dynamics: Studies in car following. *Operations Research*, 6:165–184, 1958.

[16] R.M.H. Cheng, J.W. Xiao, and S. LeQuoc. Neuromorphic controller for AGV steering. In *Proceedings, IEEE International Conference on Robotics and Automation*, 1992.

[17] E.R.F.W. Crossman and H. Szostak. Man-machine models for car steering. In *Fourth Annual NASA-Univ. Conference on Manual Control*, pages 171–195, March 1968.

[18] Kan Deng. *OMEGA: ON-LINE MEMORY-BASED GENERAL PURPOSE SYSTEM CLASSIFIER*. PhD thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, November 1998.

[19] T.A. Dingus, H.L Hardee, and W.W. Wierwille. Development of models for on-board detection of driver impairment. *Accident Analysis and Prevention*, 19(4):271–283, 1987.

[20] Edmund Donges. A two-level model of driver steering behavior. *Human Factors*, 20(6):691–707, 1978.

[21] 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.

[22] Edward Fix. Modeling human performance with neural networks. In *Proceedings of the International Joint Conference on Neural Networks*, pages 247–252, 1990.

[23] T. Fujioka and N. Takubo. Driver model obtained by neural network system. *JSAE Review*, 12(2):82–85, 1991.

[24] 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.

[25] Michael A. Goodrich and Erwin R. Boer. Validatin the satisficing approach to collision avoidance system design. In *Proceedings of CESA '98*, 1998.

[26] Hideyuki Goto, Kyoko Abe, Fumio Munekata, and Kazuyuki Kobayashi. Estimation of driving loci and evaluation of driving skill. In *Proceedings of Intelligent Vehicles*, pages 388–393, 1995.

[27] Alfred Gray. *Modern Differential Geometry of Curves and Surfaces*. CRC Press, 1993.

[28] M.S. Habib. Characterization of driver-vehicle directional control using three models of human driver. In *Proceedings of the International Symposium on Advanced Vehicle Control*, pages 36–41. SAE Japan, October 1994.

[29] J.A. Hadden, J.H. Everson, D.B. Pape, V.K. Narendran, and D.A. Pomerleau. Modeling and analysis of driver/vehicle dynamics with "run-off-road" crash avoidance systems. In *Proceedings of the 30th ISATA Conference*, 1997.

[30] R. A. Hess and A. Modjtahedzadeh. A control theoretic model of driver steering behavior. *IEEE Control Systems Magazine*, pages 3–8, August 1990.

[31] Petros A. Ioannou. Autonomous intelligent cruise control. *IEEE Transactions on Vehicular Technology*, 42(4):657–672, November 1993.

[32] Kazunori Isomoto, Tadayuki Nibe, Takamasa Suetomi, and Tetsuro Butsuen. Development of a lane-keeping system for lane departure avoidance. In *Proceedings of the 2nd World Congress on Intelligent Transportation Systems*, number 2, pages 1266–1271, Yokohama, Japan, November 1995.

[33] Yutaka Kanayama and Norihisa Miyake. Trajectory generation for mobile robots. In *Proceedings of the International Symposium on Robotics Research*, volume 3, 1985.

[34] Bart Kosko. *Fuzzy Engineering*. Prentice Hall, 1997.

[35] Ronalrd 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.

[36]]Rainer Koenig, Axel Saffran, and Hans Breckle. Modelling of drivers' behavior. In *Vehicle Navigation and Informations Systems Conference Proceedings*, pages 371–376, 1994.

[37] K. F. Kraiss. Implementation of user-adaptive assistants with neural operator models. *Control Engineering Practice*, 3(2):249–256, 1995.

[38] Chris Kreucher, Sridhar Lakshmanan, and Karl Kluge. A driver warning system based on the LOIS LANE detection algorithm. In *Proceedings of the 1998 IEEE International Conference on Intelligent Vehicles*, Stuttgart, Germany, November 1998. IEEE.

[39] M.F. Land and D.N. Lee. Where we look when we steer. *Nature*, 369:742–744, June 1994.

[40] Michael Land and Julia Horwood. Which parts of the road guide steering. *Nature*, 377:339–340, September 1980.

[41] Chiu-Feng Lin. *Lane Sensing and Path Prediction for Preventing Vehicle Road-Departure Accidents*. PhD thesis, University of Michigan, 1995.

[42] Andrew Liu and Alex Pentland. Towards real-time recognition of driver intentions. In *Proceedings, IEEE Intelligent Transportation Systems Conference*, Boston, MA, November 1997. IEEE.

[43] G. Malaterre and D. Lechner. Emergency maneuvers at junctions using a driving simulator. In *Transportation and Traffic Theory: Proceedings of the 11th International Symposium on Transportation and Traffic Theory*, pages 213–232, 1990.

[44] Charles C. MacAdam and Gegory E. Johnson. Application of elementary neural networks and preview sensors for representing driver steering control behavior. *Vehicle System Dynamics*, 25:3–30, 1996.

[45] D.T. McRuer and E. Krendel. Mathematical models of human pilot behavior. In *NATO AGARDograph*, number 188, 1974.

[46] Duane T. McRuer, R. Wade Allen, David H. Weir, and Richard H. Klein. New results in driver steering control models. *Human Factors*, 19(4):381–397, 1977.

[47] Klaus Mecklenburg, Tomas Hrycej, Uwe Franke, and Hans Fritz. Neural control of autonomous vehicles. In *Proceedings of IEEE Vehicular Technology Conference*, pages 303–306, 1992.

[48] Andrew Moore. *Efficient Memory-based Learning for Robot Control*. PhD thesis, Cambridge University, March 1991.

[49] K. Naab and G. Reichart. Driver assistance systems for lateral and longitudinal vehicle guidance - heading control and adaptive cruise control. In *Proceedings of the International Symposium on Advanced Vehicle Control*, pages 449–454, 1994.

[50] Michael C. Nechyba and Yangsheng Xu. Stochastic similarity for validating human control strategy models. In *Proceedings of the International Conference on Robotics and Automation*, 1997.

[51] Michael C. Nechyba and Yangsheg Xu. On discontinuous human control strategies. In *Submitted, Proceedings IEEE Int. Conference on Robotics and Automation*. IEEE, 1998.

[52] S. Neusser, B. Hoefflinger, and J. Nijhuis. A case study in car control by neural networks. In *Proceedings of the ISATA International Symposium*, Florence, Italy, 1991. INRETS.

[53] NHTSA. The Autonav/DOT project: Baseline measurement system for evaluation of roadway departure warning systems. Technical Report DOT HS 808 895, U.S. Department of Transportation: National Highway Traffic Safety Administration, http://www.nhtsa.dot.gov, January 1999.

[54] R. Onken and J.P. Feraric. Adaptation to the driver as part of a driver monitoring and warning system. In *The Second International Conference on Fatigue and Transportation: Engineering, Enforcement, and Education Solutions.*, February 1996.

[55] Doug Pape, Jeff Hadden, V.K. Narendran, Nathan Brown, Jeff Everson, Dean Pomerleau, and Charles Thorpe. Effectiveness analysis of run-off-road countermeasures for passenger vehicles on freeways and two-lane secondary roads. Available from US Department of Transportation, National Highway Traffic Safety Administration, Office of Crash Avoidance Research, Washington, DC, 20590.

[56] Doug Pape, Jeff Hadden, V.K. Narendran, Nathan Brown, Jeff Everson, Dean Pomerleau, and Charles Thorpe. Effectiveness analysis of run-off-road countermeasures for heavy trucks. Available from US Department of Transportation, National Highway Traffic Safety Administration, Office of Crash Avoidance Research, Washington, DC, 20590.

[57] D.B. Pape, J.A. Hadden, N.J. McMillan, V.K. Narendran, J.H. Everson, and D.A. Pomerleau. Performance considerations for run-off-road countermeasure systems for cars and trucks. In *SAE Technical Paper*, number 1999-01-820, pages 113–118, March 199.

[58] Thomas E. Pilutti. *Lateral Vehicle Co-Pilot To Avoid Unintended Roadway Departure*. PhD thesis, The University Of Michigan, 1997.

[59] L.A. Pipes. An operational analysis of traffic dynamics. *Journal of Applied Physics*, 24:271–281, 1953.

[60] Dean A. Pomerleau. Personal communication.

[61] Dean Pomerleau. *Neural network perception for mobile robot guidance*. Kluwer Academic Publishing, 1993.

[62] Dean A. Pomerleau. RALPH: Rapidly adaptive lateral position handler. In *Proceedings, IEEE Symposium on Intelligent Vehicles*, Detroit, MI., Sept. 1995. IEEE.

[63] D.A. Redelmeier and R.J. Tibshirani. Association between cellular-telephone calls and motor vehicle collisions. *New England Journal of Medicine*, 1997

[64] Douglas A. Reece and Steven A. Shafer. A computational model of driving for autonomous vehicles. *Transportation Research*, 27A(1):23–50, 1993.

[65] Josef Schumann, Jan Loewenau, and Karl Naab. The active steering wheel as a continuous support for the driver's lateral control task. In *Proceedings of Vision in Vehicles V*, pages 1–8, 1995.

[66] J.F. Shepanski and S.A. Macy. Manual training techniques of autonomous systems based on artificial neural networks. In *Proceedings of the IEEE First Int'l Conference on Neural Networks*, pages 697–704, 1987.

[67] Gunter P. Siegmund, David J. King, and David K. Mumford. Correlation of heavy-truck driver fatigue with vehicle-based control measures. *SAE Transactions: Commercial Vehicles*, 3(SAE 952594):441–468, 1995.

[68] Don Stauffer and James Lenz. An electronic rumble strip. *SPIE International Society of Optical Engineering*, 2902:106–112, 1997.

[69] National Highway Transportation and Safety Administration. Fatality analysis reporting system. On the web at http://www-fars.nhtsa.dot.gov/fars.

[70] Hiroshi Takahashi and Kouichi Kuroda. A study on automated shifting and shift timing using and driver's mental model. In *Proceedings, Intelligent Vehicles*, 1996.

[71] Louis Tijerina, James L. Jackson, Dean A. Pomerleau, Richard A. Romano, and Andrew D. Petersen. Driving simulator tests of lane departure collision avoidance systems. In *Proceedings of ITS America Sixth Annual Meeting*. ITS America, April 1996.

[72] R. Tribe, K. Prynne, and I. Westwood. Intelligent driver support. In *Proceedings of the 2nd World Congress on Intelligent Transportation Systems*, pages 1187–1192, Yokohama, Japan, Nov 1995.

[73] Walter W. Wierwille, Gilbert A. Gagne, and James R. Knight. An experimental study of human operator models and closed-loop analysis methods for high-speed automobile driving. *IEEE Transactions on Human Factors in Electronics*, 8(3):187–201, 1967.

[74] D.H. Weir and D.T. McRuer. Measurement and interpretation of driver/vehicle system dynamic response. *Human Factors*, 15:367–378, 1973.

[75] Neal E. Wood. Shoulder rumble strips: A method to alert drifting drivers. Pennsylvania Turnpike Commission Report, 1994.

[76] Liang Zhao and Charles Thorpe. Qualitative and quantitative car tracking from a range image sequence. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 496–501. IEEE, June 1998.