

Final Report

Run-Off-Road Collision Avoidance Countermeasures Using IVHS Countermeasures

TASK 3-Volume 1

NOTE TO READER:

THIS IS A LARGE DOCUMENT

Due to its large size, this document has been segmented into multiple files. All files separate from this main document file are accessible from links ([blue type](#)) in the [table of contents](#) or the body of the document.



U.S. Department
of Transportation
**National Highway
Traffic Safety
Administration**



DOT HS 808 501

August 1995

Final Report

Run-Off-Road Collision Avoidance Countermeasures Using IVHS Countermeasures

TASK 3-Volume 1

This publication is distributed by the U.S. Department of Transportation, National Highway Traffic Safety Administration, in the interest of information exchange. The opinions, findings and conclusions expressed in this publication are those of the author(s) and not necessarily those of the Department of Transportation or the National Highway Traffic Safety Administration. The United States Government assumes no liability for its contents or use thereof. If trade or manufacturers' name or products are mentioned, it is because they are considered essential to the object of the publication and should not be construed as an endorsement. The United States Government does not endorse products or manufacturers.

Contracting Officer's Technical Representative's Precis

This report provides a basis for disseminating the preliminary contract results on a timely basis resulting in the information being available before the contract final reports are produced. Research performed during the remainder of the contract may support and/or modify the results, therefore, the material contained in this report should not be consider to be final. The current schedule calls for the completion of this research project by the third quarter of 1999.

1. Report No DOT HS 808 501	2. Government Accession No	3. Recipients's Catalog No	
4. Title and Subtitle Run-Off-Road Collision Avoidance Countermeasures Using IVHS Countermeasures Task 3 Report - Volume 1		5. Report Date August 23, 1995	
		6. Performing Organization Code	
7. Author(s) D. Pomerleau, P. Kumar, J. Everson, L. Lazofson, E. Kopala		8. Performing Organization Report No	
9. Performing Organization Name and Address Robotics Institute Carnegie Mellon University 5000 Forbes Avenue Pittsburgh, PA 15213		10. Work Unit No (TRAIS) n code	
		11. Contract or Grant No DTNH22-93-C-07023	
12. Sponsoring Agency Name and Address National Highway Traffic Safety Administration 400 Seventh Street, S.W. Washington, DC 20590		13. Type of Report and Period Covered Final Report 10-93 to 8-95	
		14. Sponsoring Agency Code	
15. Supplementary Notes			
16. Abstract The Run-Off-Road Collision Avoidance Using IVHS Countermeasures program is to address the single vehicle crash problem through application of technology to prevent and/or reduce the severity of these crashes. This report describes the findings of the Task 3 effort. Task 3 focused on testing of existing technology to meet the functional goals for run-off-road countermeasures identified in Task 2. These tests included all aspects of countermeasure performance, including sensing functions, algorithm or decision making functions and driver interface functions. Tests of existing technology were performed using a range of techniques and facilities, including laboratory experiments, in-vehicle tests and driving simulator tests. Two primary categories of run-off-road countermeasure technologies were tested in this effort - lateral countermeasures and longitudinal countermeasures. Lateral countermeasures are designed to prevent run-off-road crashes in which the vehicle drifts from its lane because of driver inattention or because the driver relinquishes steering control due to drowsiness, intoxication or some other medical condition. Technology tested in this category included forward and downward looking vision systems for sensing the vehicle's lateral position on the roadway. Longitudinal countermeasures are designed to prevent run-off-road crashes in which the vehicle departs the road due to excessive speed for the roadway geometry or pavement conditions. Technology tested in this category included a combination of GPS and digital maps for sensing the the distance to, and the severity of, upcoming curves.			
17. Key Words Vision Syst. GPS	Run-Off-Road Collision Avoidance Single Vehicle Roadway Departure Lateral Countermeasures Longitudinal Countermeasures	18. Distribution Statement Document is available to the public through the National Technical Information Service, Springfield, VA 22161	
19. Security Classif (of this report) Unclassified	20. Security Classif (of this page) Unclassified	21. No of Pages 113 + appendices	22. Price \$1,217,951

Foreword

The Run-Off-Road Collision Avoidance Using IVHS Countermeasures program is to address the single vehicle crash problem through application of technology to prevent, and/or reduce the severity of, these crashes. The prime contractor for this effort is Carnegie Mellon University (CMU) operating under Contract No. DTNH22-93-C-07023. Members of the project team include Battelle Memorial Institute, Calspan Corporation and the University of Iowa.

The program consists of a sequence of nine related tasks to be completed in three distinct program phases. Phase I of this effort is currently fully funded and is comprised of the first four program tasks. Primary task completion responsibility has been assigned to individual team members with Calspan conducting Tasks 1 and 2, CMU conducting Task 3, and Battelle conducting Task 4. As prime contractor, CMU provides guidance and oversight to all subcontractor efforts.

The Task 1 and Task 2 efforts have been completed. The Task 1 effort involved characterizing the circumstances in which run-off-road crashes occur by analyzing the national crash databases (NASS CDS and GES databases). The Task 2 effort involved classifying the crash circumstances into similar subsets, identifying opportunities for intervention within these subsets, and formulating functional goals for countermeasures that would prevent or reduce the severity of run-off-road crashes.

The Task 3 effort, described in this report, focused on testing of existing technology to meet the functional goals identified in Task 2. These tests included all aspects of countermeasure performance, include sensing functions, algorithm or decision making functions and driver interface functions. Tests of existing technology were performed using a range of techniques and facilities, include laboratory experiments, in-vehicle tests and driving simulator experiments.

Technical results from the Task 3 tests will be utilized in Task 4 to develop computer models of countermeasure effectiveness, and to develop preliminary specification for run-off-road countermeasure performance characteristics. In addition, it is anticipated that this report will function as a resource for reference for Phase II and Phase III tasks.

Table of Contents

1.0 Introduction	1
2.0 Approach	4
2.1 Identify technology to be tested	4
2.2 Acquire/Build Technology	7
2.3 Design Tests	7
2.4 Conduct and Document Tests	7
3.0 Lateral Countermeasure Sensing/Algorithm Tests	9
3.1 Characteristics of Laterally Induced Crashes	9
3.2 Functional Goals for Lateral Countermeasures	10
3.3 Goal 1: Monitor vehicle dynamic state	10
3.4 Goal 2: Determine Vehicle Position/Orientation Relative to Road	10
3.4.1 Infrastructure-based Lateral Position Detection Systems	11
3.4.2 Lateral Position Detection Systems without Forward Preview	11
3.4.3 Lateral Position Detection Systems with Forward Preview	23
3.5 Goal 3: Inferring Driver's Intentions	57
3.6 Goal 4: Detect Potential for Roadway Departure	58
3.6.1 Time-to-Line-Crossing (TLC) Algorithm	58
3.6.2 Time-to-Trajectory-Divergence (TTD) Algorithm	60
3.7 summary	62
4.0 Longitudinal Countermeasure Sensing/Algorithm Tests	64
4.1 Functional Goals	65
4.2 Goal 1: Monitoring Vehicle Dynamic Status	66
4.2.1 Vehicle Velocity	66
4.2.2 Vehicle Acceleration/Deceleration	66
4.2.3 Implementation and Test Results	67
4.3 Goal 2: Determine upcoming Road/Curve Geometry	67
4.3.1 Direct Measurement	67
4.3.2 Transponders	67
4.3.3 Commercial Map Databases	68
4.3.4 Custom Built Maps	68
4.3.5 Implementation and Test Results	69
4.4 Goal 3: Determine Vehicle Longitudinal Position Relative to Curve	75
4.4.1 Direct Measurement	75
4.4.2 Transponders	75
4.4.3 GPS/DGPS based vehicle location	76
4.4.4 Implications of Results	88
4.5 Goal 4: Detect Degraded Roadway Conditions	88
4.5.1 Infrastructure-Based Sensing of Roadway Conditions	89
4.5.2 In-Vehicle Sensing of Roadway Conditions	91
4.5.3 Determining the Coefficient of Friction	92
4.5.4 Implementation and Testing	93
4.6 Goal 5: Process data to determine acceptable speed for upcoming road	93
4.6.1 Safe Speed Estimation	93
4.6.2 Integrated Longitudinal Sensing and Processing Algorithm	94

4.7 Goal 6:Present phased alarm to driver	96
4.8 Results of Integrated tests	97
4.8.1 Repeatability of Curve Alert Warnings	98
4.9 Summary	99
5.0 Summary	100
5.1 Lateral Technology Tests	100
5.2 Longitudinal Technology Tests	101
5.3 Driver Interface Tests	102
5.4 Conclusions	102
References	103
Appendix A: Description of Testbed Vehicle	106
Appendix B: Sony 711 Camera Calibration Data	113

List of Tables

Table 3- 1: Lateral displacement estimates from Aurora and manual measurement	22
Table 3-2: Status of forward-looking lateral position systems	24
Table 3-3: RALPH lane location estimation accuracy	46
Table 4-1: Roadway alignment: fatal vs. all run-off-road crashes.....	64
Table 4-2: Violations charged by horizontal alignment	64
Table 4-3: Roadway alignment in SVRD crashes - CDS data (weighted %).	65
Table 4-4: Causal factor by horizontal alignment.....	65
Table 4-5: Current and anticipated capabilities of GPS receivers	77
Table 4-6: GPS satellites tracking statistics..	83
Table 4-7: Lateral friction coefficient for various road/tire conditions	92
Table 4-8: Variability of longitudinal countermeasure warning onset time	99

List of Figures

Figure 2- 1: Block diagram of run-off-road countermeasure functions..	5
Figure 3- 1: Downward looking roadway departure warning system.....	13
Figure 3-2: A typical image of a lane marker on the road	13
Figure 3-3: Real image with overlay of the scanline intensity profile	14
Figure 3-4: Basic shape of the template used for correlation in lane marker tracking..	15
Figure 3-5: Adjustable template..	16
Figure 3-6: Illustration of the adjustable template as a function of pixel position..	16
Figure 3-7: Scanline intensity profile, contrast, and match error	17
Figure 3-8: Comparison of step edge template and the template with gaps	18
Figure 3-9: Camera calibration using marks evenly spaced at known intervals..	19
Figure 3-10: Display on the monitor once a lane marker is detected	20
Figure 3-1 1: Template for double yellow lane marker	21
Figure 3- 12: Display on the monitor once a double lane marker is detected	21
Figure 3-13: AURORA's estimate of lateral displacement over time..	22
Figure 3- 14: ALVINN camera view and preprocessed image..	26
Figure 3- 15: ALVINN neural network architecture	27
Figure 3-16: ALVINN image transformation scheme..	28
Figure 3-17: Color bands of rural road without lane markings at various visibilities..	33
Figure 3-18: Color bands of rural road with yellow centerline at various visibilities..	34
Figure 3- 19: Color bands of multi-lane divided highway at various visibilities	35
Figure 3-20: Color images of three road types degraded at various visibilities	36
Figure 3-21: Mean trajectory divergence as function of visibility for three road types..	38
Figure 3-22: Standard deviation of trajectory divergence vs. visibility for three roads..	39
Figure 3-23: Forward looking image (left), and RALPH's sampling strategy (right).....	41
Figure 3-24: RALPH curvature hypotheses	42
Figure 3-25: RALPH curvature scoring technique..	43
Figure 3-26: RALPH lateral offset determination technique	44
Figure 3-27: RALPH processing a daytime highway image	47
Figure 3-28: RALPH processing a daytime highway image with heavy shadows.....	47
Figure 3-29: RALPH processing a nighttime highway image..	48
Figure 3-30: RALPH processing a daytime rural road image	48
Figure 3-3 1: RALPH processing early morning rural road image with glare off road.....	49
Figure 3-32: RALPH processing a nighttime rural road image..	49
Figure 3-33: S-curve used for testing RALPH	51
Figure 3-34: RALPH's curvature estimate on two traversals through the s-curve.....	51
Figure 3-35: Lane deviation in normal driving, and when the driver is distracted..	52
Figure 3-36: Examples of well marked roadway encountered in cross country test..	54
Figure 3-37: Roads without strong markings (left) and with wet pavement (right).	55
Figure 3-38: Road with severely worn markings (left) and unpaved road (right).	55
Figure 3-39: California freeways with reflectors instead of painted lane markings.....	56
Figure 3-40: Challenging images from city driving	56
Figure 3-41: AURORA's estimate of TLC before lane crossing	59
Figure 3-42: Derivation of the Time-to-Trajectory-Divergence (TTD) algorithm.....	61
Figure 4- 1: Sample Etak map data..	68

Figure 4-2: Moving map display system	70
Figure 4-3: Etak map of 100 km test run	71
Figure 4-4: Distribution of discrepancies between Etak map and DGPS map..	72
Figure 4-5: Curvature data extracted from Etak and custom map databases..	73
Figure 4-6: Histogram of curvature difference between Etak map and recorded map..	74
Figure 4-7: Histogram of lengths of Etak road segments	75
Figure 4-8: Position estimates from stationary GPS receiver..	78
Figure 4-9: Position estimates from stationary GPS receiver with Omnistar DGPS.....	79
Figure 4-10: Position estimates from stationary GPS receiver with Navstar DGPS	80
Figure 4- 11: Data from long baseline DGPS test	81
Figure 4-12: Detailed data from long baseline DGPS test.....	82
Figure 4-13: Data from GPS latency experiment	83
Figure 4- 14: Histogram of difference between curvatures from Etak recorded maps	84
Figure 4-15: Position data near downtown area	85
Figure 4-16: Position data near a sharp curve..	85
Figure 4-17: Position data near a cloverleaf..	86
Figure 4-18: Position data on a straight road segment..	86
Figure 4-19: Position data near a shallow curve..	87
Figure 4-20: Data collected during repeated traversals of a divided highway	87
Figure 4-2 1: Data collected during repeated U-turns at divided highway exit..	88
Figure 4-22: Data from the SCAN pavement monitoring system	90
Figure 4-23: Excessive speed warning system algorithm.....	95
Figure 4-24: User interface for the longitudinal countermeasure..	97
Figure 4-25: Longitudinal countermeasure block diagram.....	98

1.0 Introduction

Run-off-road crashes, also called Single-Vehicle Roadway Departure (SVRD) crashes, are defined in this program to include all one vehicle crashes where the first harmful event occurs after the vehicle left the road surface, except for backing and pedestrian related crashes. As was determined in Task 1, these crashes are caused by a variety of factors, including:

- Driver inattention - typically due to internal or external distraction
- Driver incapacitation - typically drowsiness or intoxication
- Evasive maneuvers - driver steers off road to avoid obstacle
- Lost directional control - typical due to wet or icy pavement
- Excessive speed - traveling too fast to maintain control
- Vehicle failure - typically due to tire blowout or steering system failure

Single vehicle run-off-road crashes represent the most serious crash problem within the national crash population. Analysis of the 1992 NASS GES file, conducted as part of the Task 1, indicated that approximately 1.21 million police-reported crashes of this type occurred in the US in that year. This number represented approximately 20.1 percent of the crashes in the GES database. In addition, more than 520,000 vehicle occupants were injured in run-off-road crashes in 1992 and this level of injury represented approximately 26.8 percent of the injuries in the GES database. In a similar manner, the 14,031 fatalities sustained in run-off-road crashes (FARS data) represented approximately 41.5 percent of the 33,846 in-vehicle fatalities that occurred in 1992 in the US. Thus, in terms of injury frequency and severity, run-off-road crashes are an extremely serious problem.

The goal of the Run-Off-Road Collision Avoidance Using IVHS Countermeasures program is to address this crash problem through the application of technology to prevent, or reduce the severity of, these crashes. Advances in sensor and processing capabilities over the past decade allow for real time collection and analysis of the information characterizing the vehicle's operating environment and the driver's performance. Application of these technologies is an integral part of a program intended to dramatically improve vehicle safety. This program, titled Intelligent Transportation Systems (ITS), formally titled Intelligent Vehicle Highway System (IVHS), will address the run-off-road crash problem, as well as a broad spectrum of transportation issues.

The current program consists of a sequence of nine related tasks to be completed in three distinct phases. Phase I of this effort is currently underway, and is comprised of the four tasks summarized below:

- Task 1: Thoroughly Analyze the Crash Problem
- Task 2: Establish Functional Goals
- Task 3: Conduct Hardware Testing of Existing Technologies
- Task 4: Develop Preliminary Performance Specifications Based on Critical Factors and

Models of Crash Scenarios

The Phase I work flow is linear in nature in that the output of one task is utilized as an input to subsequent tasks. In Task 1, for example, data analyses were conducted to determine the circumstances associated with run-off-road crashes and the reasons why these crashes occurred. Engineering evaluations were also completed to establish the dynamic states of involved vehicles and the sequence of events associated with the crashes. These results were carried forward to Task 2 where a taxonomy was developed to classify the run-off-road scenarios in terms of the relative length of time over which the road departure occurred. This information was used to develop practical functional goals for potential countermeasures.

In the Task 3 effort described in this report, the functional goals developed in Task 2 were used to formulate complete run-off-road countermeasures. These countermeasures were built and tested in situations that were identified in Task 1 to be representative of roadway departure crashes. The results of these testing efforts will be incorporated into the mathematical models developed for Task 4, and thereby influence the preliminary performance specifications developed in that effort.

Subsequent phases of this program will continue the development sequence. For example, in Phase II the contract team will perform state-of-the-art technology reviews and design test bed systems. The test bed systems will be evaluated in Phase III. The results of these Phase III tests will be used to modify and expand the preliminary performance specifications from Task 4.

The hardware testing for Task 3 has been completed. The purpose of this report is to describe and document the systems tested and the results obtained. Countermeasure design implications of these results are also addressed. The report is divided into two volumes. The first contains results of in-vehicle and laboratory tests of the sensor and algorithm components of run-off-road countermeasures. Volume II contains results of experiments on the Iowa driving simulator to evaluate alternative driver interfaces for roadway departure countermeasures. The format and section content for the Volume I are as follows:

Section 2: Approach

This section describes the methodology that is applied to the Task 3 test sequence. This sequence involves first identifying candidate technologies to fulfill the functional goals developed in Task 2. Next these technologies must be acquired or developed, depending on their availability. Then a sequence of tests must be designed to evaluate how well these technologies meet the individual functional goals. Finally these tests must be conducted and the results documented.

Section 3: Lateral Countermeasure Sensing/Algorithm Tests

As will be seen in the approach section, run-off-road countermeasures can be divided into two categories. The first category includes countermeasures designed to prevent roadway departure crashes caused primarily by a failure in lateral control. These crashes typically occur on straight or slightly curved sections of road, and are typically caused by driver inattention, driver incapacitation, and to some extent, lose of directional control. This sec-

tion describes technology with the potential to prevent these crashes, and the results of tests conducted to evaluate the performance of the technology. These tests focus on the sensing and decision algorithms for countermeasures designed to prevent laterally induced run-off-road crashes. Tests to evaluate the driver interface for these systems is described in Volume II of this report.

Section 4: Longitudinal Countermeasure Sensing/Algorithm Tests

The second category of run-off-road countermeasures includes systems designed to prevent roadway departure crashes caused primarily by a mistake in longitudinal control. In particular, these crashes often occur on curves, and are usually precipitated by excessive speed for the road geometry or pavement conditions. These includes crashes identified in Task 1 as being caused by excessive vehicle speed or lost directional control. This section describes technology with the potential to prevent these crashes, and the results of tests conducted on this technology to evaluate its performance. These tests focus on the sensing and decision algorithms for countermeasures designed to prevent longitudinally induced run-off-road crashes. Tests to evaluate the driver interface for these systems is described in Volume II of this report.

Section 5: Summary and Conclusions

A summary of the tests conducted for Task 3, and the results obtained is provided in this section. Implications of these results, and recommendations for further testing are also presented.

2.0 Approach

The goal of Task 3 is to test existing technology with potential to prevent run-off-road crashes. The purpose of this testing is to determine limits, boundaries and capabilities of the technology, to assist in formulating performance requirements for run-off-road collision avoidance systems. There were four steps in the effort conducted for Task 3 of this program. They were:

1. Identify technology to be tested
2. Acquire/build technology
3. Design tests of technology to evaluate performance
4. Conduct and document tests

The high level methodology associated with each of these four steps is addressed in the remainder of this section. The details of each step is provided in the following sections of the report.

2.1 Identify technology to be tested

In order to identify the hardware and software to be tested in this task, three factors were considered: the functional goals for run-off-road countermeasures developed in task 2, the efforts being conducted by other related programs, and the availability of technology for testing.

Of the three factors considered, the most important was the functional goals for run-off-road countermeasures developed in Task 2. These goals characterize the actions a run-off-road countermeasure must perform in order to be effective. The final set of 11 functional goals developed in Task 2 are:

1. Monitor vehicle dynamic status
2. Determine geometric characteristics of upcoming roadway segment
3. Determine vehicle position/orientation relative to roadway
4. Determine driver intention
5. Detect degraded roadway conditions
6. Process data to determine acceptable speed for upcoming roadway segment
7. Detect potential for roadway departure
8. Present phased alarm to driver
9. Determine driver state
10. Modulate driver control input
11. Maintain/regain safe vehicle attitude

As was discussed in the Task 2 report, only a subset of these functional goals would typically be required to prevent any particular run-off-road crash, or even any particular type of run-off-road crash. Taken together, these functional goals have the potential to eliminate a significant fraction

of the run-off-road crash population identified in Task 1. A block diagram depicting how these functional goals could be combined into an integrated run-off-road countermeasure system was presented in Figure 6-1 of the Task 2 report, and is reproduced here as Figure 2-1.




Figure 2-1: Block diagram of run-off-road countermeasure functions

There are several important aspects of the block diagram in Figure 2-1. First, the functions performed by the integrated countermeasure can be divided into three categories: sensing functions, processing (or decision algorithm) functions, and driver interface functions. Within the sensing and processing functions, there are three parallel functional sequences each leading to the issue-of an alert to the driver.

The first of these parallel functional sequences involves detecting dangerous impairment of driver state. If the driver is drowsy, intoxicated, or in some other way impaired, this sequence is intended to detect the situation and trigger a sequence of driver interface functions to prevent a crash. This functional sequence is included in the block diagram for completeness, but to avoid duplication of effort with the ongoing NHTSA driver impairment detection program [23], driver impairment detection has not be the focus of the Task 3 efforts for this program.

Instead the team's efforts have focused on testing systems for the other two functional sequences, which are termed "longitudinal" and "lateral" sequences for the purposes of this report. In the longitudinal sequence, the goal is to detect when the vehicle is traveling too fast for the upcoming roadway segment. The longitudinal sequence utilizes vehicle dynamic state and performance data in combination with information about the current pavement conditions and upcoming roadway geometry to determine the maximum safe speed for the vehicle. If the vehicle's current velocity exceeds the safe speed, a sequence of driver interface functions is triggered to alert the driver of the danger and avoid a crash. The longitudinal functional sequence is designed to prevent those run-off-road crashes caused by excessive speed and lost directional control.

The lateral functional sequence is designed to detect when the vehicle begins to depart the road. It utilizes data about the dynamic state of the vehicle, in combination with information about the geometry of the road ahead to determine if the vehicle's current position and orientation will likely lead to a roadway departure. If the likelihood of departure exceeds a threshold, a sequence of driver interface functions is triggered to alert the driver of the danger and avoid a crash. The lateral functional sequence is designed to prevent those run-off-road crashes caused primarily by driver inattention and driver relinquishes steering control.

It is important to note that two of the original six run-off-road crash causal factors identified in Task 1 are not addressed by these functional sequences. The first is crashes caused by evasive maneuvers in which the driver intentionally swerves to avoid an obstacle in the roadway, resulting in a roadway departure crash. As was indicated in the Task 1 and 2 reports, countermeasures for preventing this type of crash are currently being investigated in the NHTSA rear end collision countermeasures specifications program, being conducted by Frontier Engineering. Therefore, crashes caused by evasive maneuvers have been eliminated from consideration in the Task 3 efforts for this program.

The second crash type not addressed by the functional sequences in Figure 2-1 are crashes caused by vehicle failures. These crashes typically result from tire blowouts or loss of power steering due to engine failure. The Task 1 analysis conducted from this program indicates that crashes from these causes are relatively rare (only 3.6 percent of the run-off-road crash population). In addition, countermeasures to prevent these crashes would require redesigning automotive components in a way that is beyond the scope of this program. For these reasons, crashes caused by vehicle failure

have been eliminated from consideration in Task 3.

2.2 Acquire/Build Technology

After identifying lateral and longitudinal run-off-road functions as the ones to be investigated, the next step in Task 3 was to acquire and/or build technology to perform these functions. The teams efforts to obtain technology for lateral and longitudinal countermeasures for testing will be discussed in more detail in the subsequent sections. However it should be noted here that the technology search conducted for this task was unable to identify any existing complete countermeasure systems for either lateral or longitudinal run-off-road crashes which were available for testing.

After consulting with NHTSA on this issue, it was decided the Task 3 efforts to obtain run-off-road countermeasure technology for testing would consist of three parts. First, the commercially available components which could form part of a run-off-road countermeasure would be acquired. Second, those components necessary for a run-off-road countermeasure, but not commercially available, would be developed by the project team within the resource constraints imposed by the program. Third, integration of the component technologies into complete countermeasure systems for both lateral and longitudinal countermeasures would be carried out by the project team, again within the resource constraints imposed by the program.

2.3 Design Tests

The test strategy developed for Task 3 included experiments to evaluate all aspects of the performance of the countermeasure technology, including performance of the sensing components, decision algorithms and driver interface. As is apparent from Figure 2- 1, the sensing and decision or processing functions are tightly coupled and hence much of the testing conducted for this task evaluated the combination of sensing and decision algorithms. The sensing and decision algorithm tests conducted for lateral countermeasures is described in Section 3. The sensing and decision algorithm tests conducted for longitudinal countermeasures is described in Section 4.

In contrast to the tightly coupled interaction between sensing and decision making functions, the driver interface functions are relatively independent. Once the decision is made by the countermeasure to trigger a response, the response can proceed with little input from the sensing and decision making functions. Because of this independence, and because of the need for carefully controlled tests using naive subjects, it was decided to conduct the driver interface tests separately on the Iowa Driving Simulator (IDS). These simulator experiments were designed to investigate the relative performance of several combinations of audible and haptic feedback to the driver in order to prevent both lateral and longitudinal roadway departures. Details about these experiments and the results obtained are provided in Volume II of this report.

2.4 Conduct and Document Tests

The final step in Task 3 was to conduct and document the tests of run-off-road countermeasure

technologies. In conducting these tests there were three constraints which had to be met, requiring that tests be conducted in a variety of circumstances.

The first constraint was that tests must be conducted in as realistic conditions as possible. This constraint required that a mobile testbed be developed to allow for in-vehicle data collection. The mobile testbed developed for this effort, called Navlab 5, is a Pontiac Transport minivan equipped with sensors and processing hardware for both lateral and longitudinal data collection experiments. The Navlab 5 is described in Appendix A.

The second constraint was that tests be conducted to quantitatively characterize the performance of the countermeasure technologies. While much of this could be done by collecting data in the mobile testbed on normal roads, to achieve high levels of precision and repeatability required tests be conducted both in the laboratory and on restricted test tracks. The test facility utilized for some of these experiments was the track at the Vehicle Research and Test Center (VRTC) in East Liberty, Ohio.

Finally, testing of the driver interface components of run-off-road countermeasures required experiments in which human subjects were exposed to near roadway departure situations. For obvious safety reasons, these tests could not be conducted in a real vehicle, so the University of Iowa driving simulator was utilized for the driver interface tests.

3.0 Lateral Countermeasure Sensing/Algorithm Tests

In the Task 1 analysis conducted for this program, it was determined that a significant portion of single vehicle roadway departure crashes are caused by a failure of the driver to maintain proper lateral control of the vehicle. These crashes typically occur because of driver inattention (12.7 percent) or because the driver relinquishes steering control (20.1 percent) due to drowsiness, intoxication or some other medical condition. The lateral countermeasures tested in Task 3 of this program were designed to address these types of crashes. Models of how effective these countermeasures would be at preventing or reducing the severity to these crash types are under development in Task 4 of the program, and will be investigated further in later program phases.

This section is divided into 6 subsections. In 3.1, results from Task 1 are used to further characterize the circumstances surrounding single vehicle roadway crashes that result from a failure of the driver to maintain lateral control. Section 3.2 outlines the functional goals a lateral countermeasure should achieve in order to prevent or reduce the severity of these crashes. These functional goals are based on the results of the Task 2 analysis conducted for the program. Sections 3.3 through 3.6 form the heart of this chapter. They describe the results of tests performed on technology to achieve the functional goals for a lateral countermeasure. Finally section 3.6 summarizes the results of the lateral countermeasure tests, and provides recommendations for further testing.

3.1 Characteristics of Laterally Induced Crashes

The clinical analysis of 200 NASS cases conducted as part of Task 1 provides a wealth of information characterizing the circumstances surrounding roadway departure crashes. For instance, approximately 75 percent of driver inattention or driver relinquishing steering control crashes involve the vehicle departing off the right edge of the roadway. This makes intuitive sense, due to the fact that on undivided roadways where the majority of SVRD crashes take place, an excursion towards the left, into the opposing lane, provides the driver with more time to recover when there is no opposing traffic. When there is opposing traffic, an excursion towards the left can result in a head-on crash, which is not considered part of the single vehicle roadway departure category.

The roadway alignment in these crash situations varies substantially - 65.7 percent of SVRD crashes caused by driver inattention happen on curves, while only 37.5 percent of the crashes involving driver relinquishing steering control occur on curves. This can be explained by the fact that periods of inattention are typically quite brief, and therefore typically only result in crashes when precise steering maneuvers are critical, such as when negotiating a curve. The duration of steering failure on the part of the driver are presumably longer when the driver is incapacitated, and therefore many more SVRD crashes caused by the driver relinquishing steering control happen on straight sections of road.

Interestingly, the vast majority of both driver inattention crashes and driver relinquishes steering control crashes occur when there are no adverse weather conditions. For driver inattention crashes, 100 percent of the sampled NASS cases occurred under favorable weather conditions. For relinquishes control crashes the corresponding figure was 86.4 percent.

The engineering analysis conducted for Task 1 indicates road departure angles in these crash categories was relatively shallow, typified by the vehicle slowly drifting off the road. Despite the low departure angle in these situations, in only 10.5 percent of the crashes does the driver attempt a corrective action to avoid the crash while the vehicle is still on the road.

Together these findings suggest that a countermeasure which can detect when the vehicle is about to depart from the road and triggers either a warning or some form of control intervention could potentially prevent many of these crashes.

3.2 Functional Goals for Lateral Countermeasures

In order to better specify the actions to be performed by a lateral roadway departure countermeasure, a list of functional goals appropriate for preventing this type of crash was developed in Task 2. These functional goals include:

1. Monitor vehicle dynamic state
2. Determine vehicle's position/orientation relative to road
3. Infer driver's intentions
4. Detect potential for roadway departure
5. Present phased warning to driver

The remainder of this section focuses on tests of technology for identifying situations in which substantial danger of laterally induced roadway departure crash exists. In doing so, this section focuses primarily on achieving goals 2 and 4. The analysis and assessment of technology for achieving goal 1, determining vehicle dynamic state, is addressed in Section 4.2 on longitudinal countermeasures, and in Appendix A. A discussion of Goal 3 is included in this section, although the project team could identify no available technology for accomplishing this goal, and so tests were not conducted. Tests of the driver interface for lateral departure countermeasures are presented in the Volume II of this report.

3.3 Goal 1: Monitor vehicle dynamic state

The primary dynamic state variable required by a lateral countermeasure is vehicle velocity. The vehicle's velocity is one of the principle factors determining the time available before a roadway departure. Technology for measuring vehicle speed was assessed in tests of the longitudinal countermeasure, which are discussed in detail in Section 4.2. For the purposes of a lateral countermeasure, several technologies are available for measuring vehicle speed to the accuracy required. These include the vehicle's speedometer, and the Doppler-based velocity estimates provide by Global Position Sensors (GPS).

3.4 Goal 2: Determine Vehicle Position/Orientation Relative to Road

The first challenging function a lateral countermeasure must perform is to sense the vehicle's

position and orientation relative to the roadway. This sensing function could potentially be accomplished in a number of different ways. Technologies for monitoring the instantaneous lateral position of the vehicle can be divided into infrastructure-based technologies, which require modifications to the roadway, and vehicle-based technologies which rely on existing roadway characteristics and in-vehicle processing.

3.4.1 Infrastructure-based Lateral Position Detection Systems

Infrastructure-based lateral position detection systems typically exploit ferromagnetic markers (usually wires or magnets) buried in the pavement. In-vehicle sensors detect these signals and use their intensity to calculate lateral position. One of the most advanced system of this type has been developed and tested by the California PATH group (a consortium of Universities and the California Department of Transportation). Statistics provided by the PATH team indicate that by using magnets buried at one meter intervals on their Berkeley test track, they can achieve lateral position estimation accuracy on the order of several centimeters under a variety of conditions [3].

While this technology is capable of impressive lateral positioning accuracy, there are several technical drawbacks that limit its deployability. One technical shortcoming of ferromagnetic detection systems is that they have difficulty when other metal is embedded in the roadway, for instance on bridge decks. A more limiting shortcoming of these systems is the logistical difficulty of deployment and maintenance. To be an effective countermeasure, a lateral position detection system must work on rural roadways, since the Task 1 analysis conducted for this program indicates that over 2/3rds of all roadway departure crashes occur on rural roads. The cost of deploying and maintaining the embedded markers on all of the nations 4 million miles of rural roadways would be prohibitively expensive, particularly in northern areas of the country where the markers would most likely result in increased pothole formation. This drawback is such a concern that we have been told by a representative of the Minnesota Department of Transportation that they will not even consider countermeasures which require embedding markers in the pavement [4].

An alternative ferromagnetic lateral position estimation system currently under development by 3M Corporation [16] relies on magnetic tape which can be stuck onto the road surface or pressed into the pavement during construction. While not available for testing in Task 3, this technology warrants further evaluation since it has the potential to provide high reliability lateral position estimates with fewer deployment difficulties than buried wires or magnets.

3.4.2 Lateral Position Detection Systems without Forward Preview

An alternative to infrastructure-based systems for lateral position detection are systems mounted on the vehicle which sense characteristics of existing roadways. These systems often use downward-looking video or infrared sensors to detect the position of the road's lane markings. The project team identified two such systems for testing in Task 3.

3.4.2.1 Laser-based Downward-looking Lateral Positioning Systems

The first downward-looking system identified by the team was an infrared-based system developed by Aerometrics Inc. under NHTSA sponsorship. The Aerometrics sensor uses a scanning infrared laser pointed down at the road, mounted behind the front license plate. The laser scans nearly 180 degrees laterally across the roadway. Detectors within the device sense the intensity of laser light being reflected back to the sensor from the road surface. The concrete or asphalt reflects little of this laser light back to the detector. However the lane markers reflect large quantities of laser light back to the detector, because the lane marker paint has retroreflective glass beads embedded within it to increase visibility at night. This difference in returning laser light is used to locate the lane markers, and to estimate lateral position. The Aerometrics system provides lateral position estimates at a rate of 200-400 hz [32].

Tests performed by Aerometrics in the parking lot of their facility indicate the system is able to estimate lateral position to within one centimeter, even under wet pavement conditions. The project team had planned to perform additional tests to determine the sensitivity of the Aerometrics sensor to the condition of the lane markings. Unfortunately, due to fabrication and personnel difficulties, Aerometrics was unable to provide a unit for testing as part of Task 3.

3.4.2.2 Vision-based Downward-looking Lateral Positioning Systems

An alternative downward-looking lateral position detection system is the AURORA system developed at Carnegie Mellon University. Instead of a laser, this system uses a downward-looking video camera to detect both white and yellow lane markings. This section describes the AURORA system and the results of tests conducted by the project team to evaluate its performance. More details about the AURORA system are also provided in [4].

3.4.2.2.1 AURORA Sensor Configuration

AURORA employs a downward-looking video camera to detect lane markers alongside the vehicle. A color camera is mounted on the side of a car, pointed downwards toward the road; this enables AURORA to view an area of the road approximately 1.6m by 1.5m next to the vehicle (See Figure 3-1). The video output of the camera is captured by a digitizer and processed using a portable Sun Sparc workstation. AURORA processes both fields of every frame provided by the digitizer (a full NTSC image frame has odd and even rows which are scanned separately, resulting

in two video fields), giving it a processing rate of 60Hz.




Figure 3-1: Downward looking roadway departure warning system

3.4.2.2.2 AURORA Processing Algorithm

AURORA relies on detection of painted lane markings to determine the vehicle's lateral position. There are two common types of lane markers: double yellow lines (separating lanes of traffic travelling in opposite directions), and single dashed white lines (separating lanes of traffic travelling in the same direction). After a simple initial camera calibration, the system is able to estimate the vehicle's lateral position accurately using either type of lane marker. Figure 3-2 is a typical image taken by AURORA's camera. From this single image, AURORA outputs whether a lane marker is present in this field, plus the distance between the vehicle and the lane marker if one is present.




Figure 3-2: A typical image of a lane marker on the road

To accommodate the real-time requirement of roadway departure warning, AURORA processes

only a single scanline of each video field using a **binormalized** adjustable template correlation technique. We choose the central row scanline of each image purely for display convenience. For a symmetric neighborhood around each point along the scanline, we compute the resemblance of this neighborhood to a lane marker template. If the resemblance of the best neighborhood is above a threshold, the point under examination is defined as the position of the lane marker. If there are no neighborhoods satisfying this threshold test, AURORA indicates that there is no lane marker in the image.




Figure 3-3: Real image with overlay of the scanline intensity profile

3.4.2.2.1 AURORA Adjustable Template Correlation

A typical scanline and its corresponding intensity profile superimposed on the image are shown in Figure 3-3. AURORA models a lane marker as having an intensity profile as shown in Figure 3-4, with a uniform intensity I_{marker} distinct from the intensity of the road I_{road}

Unfortunately, real lane markers on roads differ substantially from this simple template. They typically are not solid lines: their edges are often obscured, and they are often faded and worn due to traffic and weather. Moreover, the pavement itself is neither uniform nor clean. The roadway often contains small patches of white paint, or even faint older lane markers overlapped with newly painted ones. Furthermore, different lighting and weather conditions or shadows on the road will make the lane markers appear different from one another.

Because of these imperfections, a fixed template is not sufficient to model real lane markers; an adjustable template is necessary. AURORA's initial approach was to dynamically update the template based on recent road conditions. For instance, it adapted the marker intensity I_{marker} to match




Figure 3-4: Basic shape of the template used for correlation in lane marker tracking

that of recently detected lane markers, and adapted the road intensity I_{road} according to the road condition in the previous images. While this method worked well for slowly-varying roads, the system was easily confused by sudden changes in the road surface, such as changing from old pavement to new pavement. Since the template was updated each time AURORA finds a lane marker, this method was vulnerable to severe error propagation. For example, if a noisy patch of the road was incorrectly labelled as a lane marker, the template would have adapted according to the erroneous I_{marker} and I_{road} . This could completely corrupt the template, and make automatic recovery very difficult.

Since previous images may not be accurate predictors of current conditions, AURORA uses an adjustable template method which focuses on the current image. The overall shape of the adjustable template remains fixed (like in Figure 3-4), but now I_{marker} and I_{road} are functions of the position along the scanline. At each point under examination, AURORA projects the template outline symmetrically onto its neighborhood. Then it adjusts I_{marker} and I_{road} to the average of the intensities of scanline pixels in the corresponding region of the template, as shown in Figure 3-5.

The equation below shows the process of computing I_{marker} and I_{road} in detail. $I_{scanline}$ represents the intensity of points on the scanline. R_{marker} and R_{road} are the marker region and road region on the scanline correspond to the template regions respectively.

$$I_{marker} = \frac{1}{N_m} \cdot \sum_{i \in R_{marker}(p)} I_{scanline}(i), \quad N_m \text{ is size of } R_{marker}$$

$$I_{road} = \frac{1}{N_r} \cdot \sum_{i \in R_{road}(p)} I_{scanline}(i), \quad N_r \text{ is size of } R_{road}$$

The adjustable template will resemble a dual step function only when it is applied to the area around the true lane marker, as shown at point 2 in Figure 3-6. In uniform parts of the image away from the scanline, the adjustable template will appear as a straight line, as shown at point 1 in Figure 3-6.




Figure 3-5: Adjustable template



Figure 3-6: Illustration of the adjustable template as a function of pixel position

Binormalized correlation is applied to locate the part of the scanline which most closely resembles a lane marker. First the sum of the absolute difference between the scanline intensity $I_{scanline}$ and the corresponding template value is averaged. Then this result is divided by the difference between the I_{marker} and I_{road} of the template, as is shown in following equation.

$$Error = \frac{\sum_{i \in R_{marker}} |I_{scanline}(i) - I_{marker}| + \sum_{i \in R_{road}} |I_{scanline}(i) - I_{road}|}{(N_m + N_r) \times (I_{marker} - I_{road})}$$

This division by the contrast between the marker and road regions penalizes uniform regions of the scanline, where the adapted template might otherwise match quite closely, and favors high contrast areas typical of lane markers. To further ensure that lane markers are not detected in uniform regions, a minimum contrast between I_{marker} and I_{road} is enforced. If the contrast is below this lower bound after template adjustment, the section of the scanline surrounding the examined point is considered not to contain a lane marker. If the error in the above equation is below a threshold, AURORA judges that there is a lane marker in the image. Figure 3-7 shows the relation




Figure 3-7: Scanline intensity profile, contrast, and match error

of the intensity profile of a scanline, the contrast between I_{marker} and I_{road} for each point along this scanline, and the corresponding binormalized correlation error. Note that the error achieves minimum at the center of the lane marker where the contrast reaches a maximum.

The above method still did not always perform satisfactorily because lane markers in the real world typically do not have clean step edges at both sides. Instead, the edges may be blurry, which can cause a large correlation error. AURORA overcomes this problem by ignoring the sections of the scanline around the edges of the lane marker; it only uses the more reliable data away from the marker edges to adjust the template and compute the **binormalized correlation**, as shown in Figure 3-8. This technique is similar to that employed by the SCARF road following system [6] to detect the boundary between the road and non-road regions on unstructured roads. Ignoring these regions of the template also makes the algorithm more tolerant to variations in the width of the lane marker. One potential problem with this procedure is that it may result in reduced lane marker localization accuracy. However experiments showed that this accuracy is well preserved.




Figure 3-8: Comparison of step edge template and the template with gaps

3.4.2.2.2 AURORA Local Search Strategy

Since vehicles typically travel at speeds greater than 25 m/sec on the highway, it is imperative that AURORA operate in real time. AURORA achieves its 60 Hz cycle rate by processing only a single scanline per image. Its efficiency is further improved by first searching in the vicinity of last detected marker position. This also helps to avoid confusion caused by spurious features on the road. The local search range in AURORA is twice the size of the template's marker region R_{marker} . If the lane marker is within this local search range, it can be detected with less effort than searching through the entire scanline. If AURORA cannot find a lane marker within the local search range because it lies outside the local search range or because there is no lane marker in this image, the system extends the search to the whole scanline.

3.4.2.2.3 AURORA Sensor Calibration

Since AURORA uses a relatively wide angle lens (54 degrees) to see a large area next to the vehicle, perspective effects and lens radial distortion are significant. Because of these effects, the width of a lane marker can vary significantly depending on its position in the image. To handle these variations a calibration procedure is performed in order to determine the marker width to expect at each point along the scanline.

The calibration procedure for AURORA involves the positioning of calibration marks at known distances from the vehicle in the camera's field of view, as shown in Figure 3-9. Typically these marks are placed 10cm apart. The user then manually indicates the columns in the image at which these marks appear. By relating distance between marks on the ground to the number of columns between them in the image, the system then computes a scale factor to convert centimeters to pixels for each column in the image. This scale factor is then used to precompute how wide (in pixels) a typical lane marker should appear when centered at each column of the image. The size of the marker region R_{marker} is then adjusted at run-time according to the position in the scanline currently being searched, to compensate for perspective and lens distortion effects. This calibration is also vital for vehicle lateral displacement estimation, in order to translate a lane marker position in pixels into a distance measurement. A deployed countermeasure based on AURORA would utilize a fixed camera at a known location and orientation, so this calibration procedure would be unnecessary.




Figure 3-9: Camera calibration using marks evenly spaced at known intervals

3.4.2.2.4 AURORA Lateral Position Estimation

Once AURORA locates the lane marker, the next step is to calculate the vehicle's lateral position. The same technique used to compensate the marker width for perspective effect and lens distortion is used here.

tion is used to compute the vehicle's distance from the edge of the lane. Specifically, the approximate lateral position of the vehicle is determined by comparing the location of the detected lane marker with the locations of the known calibration marks. Linear interpolation is used to determine the precise location of lane markers falling between the positions of two calibration marks. Since the calibration marks are closely spaced, this linear interpolation does not introduce significant error, as will be shown in the next section. AURORA defines the distance between the center of the vehicle and the center of the lane as the vehicle lateral position. This can be directly computed from the distance between the vehicle and the lane marker as long as the widths of the lane and vehicle are known.

3.4.2.2.3 AURORA Performance

Extensive tests of AURORA were performed by the project team under a variety of weather and lighting conditions on a variety of road types. These conditions included sunny, cloudy, rainy and snowy days, as well as both day and night operation. Tests were conducted on both two lane rural roads and divided highways. Tests were performed in the laboratory using road sequences collected on videotape, as well as on the Navlab 5 test vehicle.

Overall, test results were quite promising. AURORA is able to reliably track both dashed and continuous white and yellow lane markers. Quantitatively, on roads with dashed white lane markers, the system misses, on average, about 1 in every 100 lane markers, usually when the marker is severely faded or obscured. This type of mistake normally occurs on a single lane marker, and does not propagate to subsequent markers since the system only relies on the current image for its processing. Figure 3-10 is the display on the monitor once AURORA has detected a white lane marker.




Figure 3-10: Display on the monitor once a lane marker is detected

While we have been using white lane markers in the previous examples, AURORA in fact works equally well for double yellow lane marker. The only differences in the algorithm are that AURORA uses color information rather than simple intensity in the template, and the template

has a different shape, as is shown in Figure 3-11. Instead of using the summation of red (R), green (G) and blue (B) pixel values as intensity for each pixel, AURORA uses $R+G-2B$ as a simple means of highlighting pixels with a yellow hue in the image. A typical result of detecting double yellow lane marker is shown in Figure 3-12.




Figure 3-11: Template for double yellow lane marker




Figure 3-12: Display on the monitor once a double lane marker is detected

AURORA's lateral position estimation accuracy was measured by comparing the result given by the system and manual measurement. Table 3-1 shows the comparison of 14 randomly selected lane markers lying at different positions on the road. The average absolute error is 0.8 cm. The standard deviation of the error in position estimation is 1.05cm.

Table 3-1: Lateral displacement estimates from Aurora and manual measurement

Sample	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Aurora	142.6	130.5	120.9	117.5	114.1	100.4	94.2	78.0	62.7	52.1	47.7	38.5	34.4	28.4
Manual	143.1	132.1	122.4	120.0	112.9	100.8	94.2	77.0	62.6	52.4	47.4	37.4	34.1	27.7
Error	-0.5	-1.6	-1.5	-2.5	1.2	-0.4	0.0	1.0	0.1	-0.3	0.3	1.1	0.3	0.7

Figure 3-13 shows a plot of AURORA's estimate of the vehicle's lateral displacement from the road center over time as the driver drifts from one side of the lane to the other. Note the smoothness of AURORA's lateral displacement estimates. Also note that the vehicle's trajectory can be characterized as periods of relatively constant velocity lateral drift, punctuated by abrupt corrective maneuvers. Flat regions indicate that marker has left the sensor's field of view.




Figure 3-13: AURORA's estimate of lateral displacement over time

In general AURORA's position estimation accuracy did not degrade with adverse weather and lighting conditions. Neither wet pavement nor a thin snow cover degraded performance significantly. However once the snow became thick enough to entirely obscure the lane markings, AURORA's accuracy fell to zero. AURORA's performance was not impaired by nighttime operation, although it did require active illumination of the pavement, in the form of a 100 watt lamp mounted next to the camera, in order to consistently track lane markers. The harsh shadow cast by

the vehicle on very sunny days occasionally proved difficult for AURORA. Under these conditions, the limited dynamic range of the camera resulting in either the shadowed region being very dark or the sunny region being extremely bright. When the lane marker transitioned between these two regions, AURORA would sometimes lose track of it. A camera with a larger dynamic range would be necessary for truly robust operation.

It was also clear from the tests performed on AURORA that tracking both left and right lane markers is necessary for reliable operation. In order to compute the vehicle's lateral position, it is necessary that a system know the lane width. With a single camera tracking the left lane marker, this is impossible to compute. In addition, on very wide roads ($> 4.5\text{m}$), the left lane marker would sometimes move out of AURORA's visual field while the vehicle was still in its lane, making it impossible to accurately compute the vehicle's lateral position. A wider field of view camera would reduce this effect, but would reduce AURORA's accuracy. A second camera tracking the right lane marker would solve these problems, and also provide a redundant source of information when lane markings are obscured or degraded.

A final drawback of downward-looking lateral position detection systems like AURORA discovered in these tests is that they are unable to detect roadway departure danger until the vehicle begins to stray from the road center. This could be problematic, since the Task 1 analysis indicates that approximately 66 percent of inattention related roadway departure crashes occur on curves, where large road curvatures coupled with obstacles like guardrails close by on the roadside, leave little time between the start of lane excursion and impact. The effectiveness of countermeasures which utilize downward-looking lateral position sensing will be carefully modeled in Task 4 to determine the significance of this effect. One approach to overcoming the lack of forward preview in downward-looking lateral position detection systems is to combine them with a digital map. This would allow the countermeasure to have at least approximate knowledge of the road geometry ahead. The team recommends investigating this combination of technologies in Phase II. As an alternative, the project team investigated a category of lateral position detection systems designed to overcome this problem: lateral position detection systems with forward preview.

3.4.3 Lateral Position Detection Systems with Forward Preview

Lateral position detection systems with forward preview are typically video-based, with a forward looking camera to detect both the current lateral position of the vehicle, and the geometry of the road ahead of the vehicle. These are by far the most actively studied type of sensing technology for single vehicle roadway departures. A partial list of references to efforts in this area include [6][7][12][20][21][22][24][25][26][27][33][35]. As part of the technology identification effort conducted for Task 3, the project team identified more than 10 groups which have been or are involved in the development of forward looking lateral position detection systems including:

- General Motors
- Ford Motor Company (Jaguar Division)
- Daimler-Benz
- The National Institute of Standards and Technology (NIST)

- Helsinki University
- Rockwell International
- University of Maryland
- Carnegie Mellon University
- IMRA America
- Toyota
- Tokyo University

Contact was made with each of these organization to determine whether they had technology which could be tested as part of this program. Unfortunately, the responses received from these organizations regarding the status of their efforts, and the potential for testing their technology, was been negative for all but two of these sources. The responses received can be divided into four categories:

Discontinued:	The effort is no longer being pursued
Proprietary:	If a system is being developed, information about it is confidential
One-of-a-kind:	The system uses expensive (more than \ \$300K custom equipment 2 and the prototype system that exists is unavailable for testing
Not Mature	In the opinion of the developer, the system is not yet ready for external testing

The responses from each of the developers listed above are categorized into one or more of these four categories in Table 3-2.

Table 3-2: Status of forward-looking lateral position systems

System Supplier	Discontinued	Proprietary	One-of-a-kind	Not Mature
General Motors	X	X		
Ford Motor Co.		X		
Daimler-Benz			X	
NIST			X	
Helsinki U.				X
U. of Maryland	X			X
IMRA America		X		X
Toyota		X		X
Tokyo U.				X

The remaining two developers, Rockwell International and Carnegie Mellon University, have

technology ready for testing by NHTSA. The Rockwell system uses a forward looking video camera and a special image processing chip to detect the lane markings on the road ahead of the vehicle. Rockwell is under contract with NHTSA to refine the system, and evaluate its performance. After consultation with both Rockwell and NHTSA personnel involved with this project, it was decided not to test the Rockwell system as part of Task 3. The reasons for this decision were: 1) the Rockwell system is still undergoing refinement, for example to determine the optimal sensor placement, and 2) Rockwell will be conducting controlled tests of their system as part of their NHTSA program.

Two forward-looking lateral position detection systems developed at Carnegie Mellon University were tested as part of Task 3. The two systems, called ALVINN and RALPH, utilize a forward-looking video camera to determine the vehicle's position on the roadway, and the geometry of the road ahead. These sensing systems themselves, and experiments conducted to evaluate them, will be presented in the next two sections.

3.4.3.1 ALVINN

The ALVINN (Autonomous Land Vehicle In a Neural Network) system is a neural network-based lateral position detection system. ALVINN uses connectionist image processing techniques to detect the location of the road ahead. It learns which image features are important for detecting the road by watching as a person drives. ALVINN was developed under sponsorship from the Advanced Research Projects Agency (ARPA) of the US Department of Defense. ALVINN was originally designed to act as the navigation system for unmanned battlefield scout vehicles. However, its ability to accurately determine the position of the road ahead makes it a good candidate sensing system for a roadway departure countermeasure.

3.4.3.1.1 ALVINN Sensor Configuration

ALVINN uses a single color video camera mounted next to the rear view mirror looking forward through the windshield at the road ahead. A typical image as seen from ALVINN's camera is shown in Figure 3-14. There are several things to note about the images processed by ALVINN. First, ALVINN does not process the entire image, but only the region within the rectangle in Figure 3-14. This is done to eliminate parts of the image that are above the horizon, or parts that represent the dashboard of the vehicle. ALVINN subsamples this rectangular region to create a low resolution image, shown in the lower left corner of Figure 3-14. It is this low resolution 30x32 pixel image that is processed by ALVINN to determine the vehicle's position and the geometry of the road ahead.




Figure 3-14: ALVINN camera view and preprocessed image

3.4.3.1.2 ALVINN Processing Algorithm

ALVINN uses an artificial neural network to process the image of the road ahead. The network is a multi-layer perceptron with three layers of neurons or “units”, as shown in Figure 3-15. The input layer consists of a single 30x32 unit “retina” onto which a video image is projected. Each of the 960 input units is fully connected to the four unit hidden layer, which is in turn fully connected to the output layer. Each of the 30 output units represents a different possible steering direction to keep the vehicle on the road. The centermost output unit represents the “travel straight ahead” condition, while units to the left and right of center represent successively sharper left and right turns.

To determine the arc the driver should currently be following in order to stay on the road, an image from the video camera is reduced to 30x32 pixels and projected onto the input layer. After propagating activation through the network, the output layer’s activation profile is translated into the optimal steering arc to follow. The steering arc recommended by the network is taken to be the center of mass of the “hill” of activation surrounding the output unit with the highest activation level. Using the center of mass of activation instead of the most active output unit to determine the optimal steering direction permits ALVINN to more accurately identify the correct direction to steer. Note that this approach combines the vehicle’s lateral position in its lane with the geome-

try of the road ahead into a single indicator of how the driver should be steering. This is different from other forward-looking lateral position estimation systems, such as the RALPH system to be discussed in the next section, which maintain these two pieces of information separately. The implications of this processing strategy for a run-off-road countermeasure will be discussed in section 3.6 on decision algorithms for roadway departure warning.




Figure 3-15: ALVINN neural network architecture

Unlike most other lateral position sensing systems, ALVINN is not programmed to detect particular features such as lane markings. Instead, the artificial neural network at the heart of ALVINN learns which steering arc is appropriate for a particular scene. The network is taught to imitate the steering response of a human driver. As a person drives, the network is trained with the back-propagation algorithm using the latest video image as input and the person's steering direction as the desired output.

To facilitate generalization to new situations, variety is added to the training set by shifting and rotating the original camera image in software to make it appear that the vehicle is situated differently relative to the road ahead. This image transformation scheme is depicted graphically in Figure 3-16. The correct steering direction for each of these transformed images is created by altering the person's steering direction for the original image to account for the altered vehicle placement. So for instance, if the person were steering straight ahead, and the image were transformed to make it appear the vehicle is off to the right side of the road, the correct steering direction for this new image would be to steer towards the left in order to bring the vehicle back to the road center. Adding these transformed patterns to the training set teaches the network how to respond when the driver has made a steering mistake, without requiring the human trainer to explicitly stray from the road center and then return.

To train ALVINN on a new road type requires the system to observe a person's steering reactions

during approximately two minutes of driving. At the end of this training period, ALVINN is ready to begin recommending steering directions, which can be compared with the driver's actual steering direction to determine the danger of roadway departure. More details of the algorithm ALVINN used to identify run-off-road situations will be provided in Section 3.6 on decision algorithms and in Volume II in the experiments on driver interfaces. For a more detailed description of the ALVINN processing algorithm, See [18].




Figure 3-16: ALVINN image transformation scheme

3.4.3.1.3 ALVINN Performance

An extremely important component of sensor processing performance is their ability to cope with degraded environmental conditions. This factor is particularly important for forward-looking video-based systems such as ALVINN, which must process images containing a wide variety of degraded conditions in order to perform reliably. Most of the tests of the ALVINN system conducted for Task 3 focused on quantifying the effects of degraded environmental conditions on sensor performance.

In order for any sensing system, including ALVINN, to effectively detect the location of the road ahead there must be sufficient contrast in the sensor input between the features delineating the road and the non-road parts of the scene. These features may include lane markings, boundaries between the road and the non-road regions, or even pavement anomalies such as oil spots down the lane center. The contrast presented by these features depends on a number of factors, including the material properties of the road and background, lighting and atmospheric conditions, and sensor performance characteristics. Poor weather and visibility conditions degrade the apparent contrast seen by the sensor and reduce the available information for subsequent processing algorithms. The contrast required to effectively locate the road ahead depends heavily on the processing algorithm utilized. As an adaptive system able to modify its processing to suit the conditions

at hand, it was expected that the ALVINN processing algorithm might be reasonably tolerant of adverse environmental conditions. However as the level of adverse ambient conditions becomes more intense, it is expected that any lateral position detection algorithm will become increasingly less accurate.

Developing a quantitative model for the effects of adverse environmental conditions for a lateral position estimation system such as ALVINN is challenging, both because it is impossible to dictate particular ambient conditions (e.g. 7mm/hr rain rate with one mile visibility fog) and because it is difficult to accurately quantify the conditions that do occur. In order to overcome these difficulties, the project team conducted controlled laboratory experiments to quantify the effects of environmental conditions on ALVINN. These experiments consisted of four parts:

1. Camera calibration
2. Image data collection
3. Controlled image degradation
4. ALVINN performance testing

Each of these parts will be discussed individually in the next four sections.

3.4.3.1.3.1 Camera Calibration

The effect ambient weather and lighting conditions will have on system performance depends heavily on the characteristics of the sensor the system employs. In order to determine the response characteristics of the video camera ALVINN uses, a series of tests were conducted on ALVINN's Sony XC-7 11 CCD color camera with Computar MIOZ-1118AMS zoom lens. In order to obtain baseline measurements of signal transfer and noise characteristics, the camera's gamma correction function and electronic shutter were switched off. The camera's automatic gain control (AGC) function was turned on, with an offset of 0dB. The signal transfer characteristics were tested using a greyscale chart, a light box, a monitor and an oscilloscope. The noise levels on the three camera output channels (red, green, blue) were also measured using the oscilloscope.

The output waveforms of the red, green and blue channels were recorded to obtain the signal transfer curves at various scene luminance levels and to characterize the AGC of the camera system. At high scene luminance levels, the camera's AGC reduces the signal transfer curve to prevent signal saturation. When the brightest scene luminance levels are reduced, the AGC attempts to increase the signal transfer curve to fully utilize its dynamic range. Graphs of the signal transfer curves are included in Appendix B.

3.4.3.1.3.2 Image Data Collection

The next step in quantifying the effects of environmental conditions of system performance was to collect several sets of road scene imagery. The team's camera equipped testbed vehicle was used to collect sequences of color images comprised of sections of multi-lane divided highway, a two-lane country road without lane markings, and a country road with a yellow centerline. Each

image was tagged with the radius of curvature steered by the driver at the moment the image was captured. The driver's radius of curvature was later used as the "optimal", or desired, response output from the ALVINN image processing system during training and testing on simulated weather-degraded versions of the images.

To accurately simulate the effects of weather degradation on these sets of images required that careful measurements be taken to characterize the conditions under which the images were collected. The team utilized radiometric measurement equipment to quantify the following environmental and system characteristics:

- terrain spectral radiance (for concrete, grass, asphalt, line markings)
- solar spectral radiance
- spot meter readings of road and grass
- camera geometry and field of view (FOV)
- frame grabber dark level and saturation level

From these ground truth measurements, the team derived the following parameters:

- terrain spectral reflectivities
- digitization correction factors
- scene-viewing geometry effects
- a conversion factor for terrain radiance to digital values

The weather for the data collection episode was mostly clear with only occasional clouds. Most of the highway scenes were imaged in bright sunlight, although some of the frames were taken with the sun behind the clouds. The highway imagery contained some shadows from bridges and other vehicles. The country road images contained segments with and without shadows from trees.

During data collection, the autoiris on the Sony XC-711 camera was disconnected and the AGC was on. This set up can accommodate several orders of magnitude in ambient irradiance levels. No image stabilization was implemented while recording the imagery, since ALVINN does not utilize a stabilized camera.

3.4.3.1.3.3 Controlled Image Degradation

Given the above image sets collected under benign, well characterized conditions, the next step was to digitally degrade these images to simulate adverse environmental effects. To accomplish this step the team utilized several tools developed previously by Battelle for analyzing and simulating battlefield sensor data. These tools include the Tactical Decision Aids (TDAs) and Electro-Optical Visualization and Simulation Tool (EOVAST).

The TDAs are automated analysis tools designed to predict the performance of electro-optical, precision guided munitions and target acquisition systems as a function of target engagement geometry and environmental conditions. The TDAs have been successfully employed to predict maximum sensor performance (in terms of maximum detection, recognition, identification, lock-on, and launch ranges) in support of Desert Storm/Desert Shield, the El Dorado Canyon Mission, and operational flight tests. TDAs have been developed for long-wave and middle-wave infrared sensors, passive daylight and low-light-level television cameras, active (laser-illuminated) television systems, 1.06m nonimaging laser receivers/designators, direct view devices (telescopes, binoculars, etc.) and night vision goggles. The TDA models implement system-level and detailed component-level EO/IR models and have been extensively validated.

Electra-Optical Visualization and Simulation Tool (EOVAST), originally developed for military targeting applications, generates and displays images as they would appear to a combat crew during an actual mission, considering the sensor system and environmental conditions. EOVAST's predicted images incorporate faceted representations of targets, predicted radiometric target and background signatures generated by a thermal signature model, the degradation in contrast due to atmospheric attenuation for the modeled environment, and blurring effects for the implemented sensor system. EOVAST performs all of this rapidly; no fractals are used. The software has been designed to execute on a Sun SPARCstation using the X-Windows environment, yet can be readily tailored for operating on any host system.

The TDA and EOVAST software tools, originally developed for tactical military applications, are also suited for electro-optical sensor system analyses regarding ITS applications and were accordingly applied to the roadway departure modeling effort. The TDA/EOVAST analysis process employed for this effort utilized the images described above of actual road scenes collected under clear conditions. The effects of adverse weather and poor visibility conditions were introduced into the frames using the TDA/EOVAST software to generate simulated images with degraded contrast. Weather and illumination variables affecting sensor performance include visibility, degree of overcast, time of day and rain rate. Illumination is a derived parameter comprised of the overcast condition and time of day factors.

Prior to operating the TDA/EOVAST software on the road scene images, The team ran the MODTRAN software package to calculate radiance and atmospheric transmission parameters required as inputs to the TDA/EOVAST software. Inputs to MODTRAN consist of environmental factors and the imaging scenario geometry including sun angle (azimuth and elevation), sensor look angle, atmospheric aerosol content (i.e., fog levels), aerosol scattering phase function (i.e., Mie or Henyey-Greenstein) and standard atmosphere type and profile. The U.S. Standard Atmosphere was selected with its accompanying temperature, pressure, humidity and other climatological parameters. MODTRAN was run over the spectral range of 300 nm through 1.1 μ m to cover the performance realm of CCD camera operation. The sun angle parameter was set to zenith (directly overhead). MODTRAN outputs included calculated solar spectral irradiance ($\text{W/m}^2\text{-m}$), path scattered radiance ($\text{W/m}^2\text{-sr}$) and atmospheric transmission (unitless).

Employing the atmospheric parameters calculated by MODTRAN, the TDA/EOVAST software was used to process the road scene frames and generate artificially degraded images under different weather conditions. The software applied to this effort consisted of a new module developed

by the team to tailor the TDA/EOVAST routines for processing road scene imagery.

The image data sets generated using this software included 30 images from each of the three road types, degraded to represent visibility conditions of 700m, 300m and 100m. Note that by definition, the term “visibility” refers to the horizontal distance for which the contrast transmission of the atmosphere in daylight is two percent. The degraded images modeled the atmospheric effects of transmission loss and the addition of path radiance. These model effects are range dependent and were correctly simulated across the entire scene field of view. Together with the original images for each road type, estimated to have a visibility of approximately 1000m, the image database for testing included a total of 360 color images:

$$[3 \text{ road types}] \times [30 \text{ images per road type}] \times [4 \text{ atmospheric conditions}]$$

Examples of the degraded imagery generated are shown in Figures 3-17 through 3-20. Figures 3-17 through 3-19 depict the red, green and blue bands of a single image from each of the three road types, degraded to the four visibility levels. As can be seen from these figures, there are slight differences in contrast between each spectral band, with the blue band providing the best contrast. Figure 3-20 shows a composite color image of the same three road images, degraded to four visibility levels. Note that in the bottom row of all these figures, representing 100m visibility, the road features are nearly indistinguishable from the background.




Figure 3-17: Color bands of rural road without lane markings at various visibilities




Figure 3-18: Color bands of rural road with yellow centerline at various visibilities




Figure 3-19: Color bands of multi-lane divided highway at various visibilities

ORIGINAL

700 meters

300 meters

100 meters




Figure 3-20: Color images of three road types degraded at various visibilities

3.4.3.1.3.4 ALVINN Performance Testing

Using the sets of simulated weather-degraded road scene images, the team assessed the sensitivity of the ALVINN road-following algorithm. This analysis ultimately served to characterize the operational envelope of the sensor/algorithm performance over a range of visibility levels.

For each set of weather-degraded images (i.e., different visibility levels), the original (undegraded) images were used to generate a training set for ALVINN. This process involved artificially shifting and rotating the first five images from each sequence 60 times to produce a total of 300 training images in which the vehicle appeared shifted and rotated relative to the roads in the original five images. These transformations took into account the geometry of both the camera and lens. For each of these transformed images, a “target” radius of curvature was computed based on the human driver’s steering radius on the original image, and the magnitude of the shift and rotation in the transformed image. This procedure for generating the training set is the same process normally used for training ALVINN (For more details, see [18]).

The resulting 300 images were then used to train the ALVINN system’s artificial neural network. Training involved repeatedly presenting the 300 images to the network and teaching the system to output the correct radius of curvature for each. This training procedure required approximately two minutes for each of the three sets of images. Recall that these three image sets represented a multi-lane divided highway, two-lane country road without lane markings, and a country road with a yellow centerline.

After training a network on undegraded images from each of the three road types, the team then tested the three networks on a disjoint set of images with various levels of degradation. The testing procedure was conducted as described below.

For each of the three road types, a set of 225 test images were generated for each of the visibility conditions by shifting and rotating the remaining 25 images using the technique described above (recall that five of the 30 original images were used to build the training set). This method of augmenting the test set was necessary since the remaining 25 original images did not contain enough variety of radii of curvature to thoroughly test ALVINN. More specifically, for most of the remaining 25 images in each set, the target radius of curvature was close to straight ahead and therefore did not exercise the networks completely. The transformed set of images showed the road at a greater variety of positions and orientations, and therefore required the network to produce a wider range of output responses.

This procedure provided five sets of 225 test images for each of the three road types. The five sets were comprised of images with visibilities of 1000m (the original image set), 700m, 400m, 300m and 100m. Just as in the training procedure, each of the test images was tagged with a target radius of curvature representing the direction the driver would steer in that situation. The network trained on the undegraded (original) divided highway images was presented with the five sets of test images depicting the divided highway in various visibility conditions. The network’s estimate of the radius of curvature for each image was compared with how the driver would steer on each of the images (i.e., the “optimal” response). The larger the difference between the networks and the driver’s responses, the bigger the error generated by the network.

To make these error measurements more intuitively meaningful, they were converted to a displacement error using the following procedure. Any difference between the radius of curvature steered by the human driver and ALVINN's radius of curvature results in two diverging trajectories; if the vehicle followed the arc indicated by ALVINN, it would follow a different path than if it followed the arc indicated by the human driver. All errors reported below were generated by measuring the distance the path dictated by ALVINN would diverge from the path dictated by the human driver if the vehicle was driven along each path for one second at the typical speed for the road type being tested (60 mph for the divided highway and 35 mph for the rural two-lane roads). This trajectory divergence distance is identical to the concept of Time-to-Trajectory Divergence, described in Section 3.6 on decision algorithms.

Figures 3-21 and 3-22 summarize the results of the weather-degraded imagery experiments. The first plot shows the mean trajectory divergence distance as a function of visibility for the three road types. The X-axis represents visibility, ranging from 100m visibility on the left to 1000m visibility on the right. The Y-axis represents the mean trajectory divergence distance (in meters) for the 225 test images under each of the visibility conditions.




Figure 3-21: Mean trajectory divergence as function of visibility for three road types




Figure 3-22: Standard deviation of trajectory divergence vs. visibility for three roads

The first important characteristic to observe from these results is that, in general, as visibility decreases, the trajectory divergence measure (error) increases. Intuitively, this makes sense: as it becomes harder to see, the network makes larger errors in estimating the correct direction to steer the vehicle.

The second important aspect of the system performance evident from the plots is that this performance degradation is not linear. In fact, the performance of the system is only slightly affected by degraded conditions in which visibility is better than 400m. A more significant impact on performance can be seen when visibility is reduced to 300m. When visibility is reduced to 100m, performance is substantially degraded, with an increase in trajectory divergence of over 100 percent in the case of the divided highway.

The third important finding involves the standard deviation of the trajectory divergence measure for the three road types as depicted in the second graph. Just as the mean trajectory divergence increases as visibility decreases, so does the standard deviation of this measure. Also note the sharp increase in standard deviation beginning in the vicinity of 300m visibility.

Two additional sets of degraded image files were generated using a different sun angle (30° rear of zenith) as a MODTRAN input. These images, processed at 100m and 300m visibility levels, resulted in poorer road-following algorithm performance compared to the 100m and 300m visibility files processed with the sun angle at zenith. Altering the sun angle changed the atmospheric backscatter effects on the visible light energy and reduced contrast in the road scenes.

In summary, this experiment assessing the effect of weather-degradation on the ALVINN lateral

position sensing system provided valuable information regarding the performance of forward-looking vision systems for roadway departure prevention. It has verified expectations that if visibility is sufficiently reduced, it will have a negative impact on sensor data processing performance. Furthermore, these experiments have quantified this effect. The findings suggest that ALVINN is relatively immune to the effects of adverse environmental conditions down to a visibility level of approximately 300m. When visibility becomes severely degraded, falling below 300m, ALVINN exhibits a relatively rapid decrease in road position estimation accuracy.

Not all aspects of adverse environmental conditions have been modeled in this experiment. For example, the effects of snow obscuring the road features and effect of specular reflection off wet pavement have not been included in these analyses. While these effects are undoubtedly important and should be addressed in Phases II and III of this program, the results from this experiment alone will provide a sound basis for the mathematical models of countermeasure performance to be developed in Task 4.

In general, tests of the ALVINN forward-looking lateral position detection system indicate that it can determine the location of the lane center one second ahead of the vehicle to an accuracy of approximately 10-25 cm. Since ALVINN does not rely on specific features (such as lane markers) being visible, but instead learns to rely on whatever features are present in the image, it can localize the lane on a variety of road types in a range of weather and lighting conditions.

However there is a downside to this flexibility. In order to adapt to a new type of road (e.g. one with a different lane marker configuration), ALVINN's neural network must be retrained. This retraining procedure takes approximately two minutes, during which ALVINN is unable to provide estimates of the lane position. This "downtime" is a problem since it means a lane departure warning system based on ALVINN would be unable to provide warnings during this period. Additionally, to retrain ALVINN requires the driver be steering correctly, an assumption that may not be valid for an impaired or inattentive driver. While limited success in overcoming these shortcomings has been achieved using a library of pretrained ALVINN neural networks to limit the need for retraining, a more comprehensive method is necessary to allow for rapid reconfiguration to changes in the driving situation.

The RALPH system, described in the next section, is an alternative lateral position detection system tested as part of Task 3. RALPH maintains ALVINN's ability to adapt to new driving conditions, but is able to perform this adaptation almost instantly, without the need for explicit training by the driver.

3.4.3.2 RALPH

The RALPH (Rapidly Adapting Lateral Position Handler) system is a forward-looking lateral position detection system developed jointly by Carnegie Mellon University and AssistWare Technology [20]. RALPH decomposes the problem of steering a vehicle into three steps, 1) sampling of the image, 2) determining the road curvature, and 3) determining the lateral offset of the vehicle relative to the lane center. The output of the later two steps are combined into a steering command, which can be compared with the human driver's current steering direction as part of a road

departure warning system, using the Time-to-Trajectory-Divergence (TTD) technique described in Section 3.6.

3.4.3.2.1 RALPH Sensor Configuration

A typical scene of the road ahead, as imaged by a video camera mounted next to the rearview mirror on our testbed vehicle, is depicted on the left of Figure 3-23. RALPH can utilize either black and white or color images, using a color-based contrast enhancement technique described in [19]. Obviously much of this image is irrelevant for the driving task (e.g. the parts of the image depicting the sky or the dashboard of the vehicle). These parts of the scene are eliminated, and only the portions of the scene inside the white trapezoid are processed. While the lower and upper boundaries of this trapezoid vary with vehicle velocity (moving further ahead of the vehicle, towards the top of the image, as vehicle speed increases), they typically project to approximately 20m and 70m ahead of the vehicle, respectively.




Figure 3-23: Forward looking image (left), and RALPH's sampling strategy (right)

The second, and perhaps more important aspect of the trapezoid's shape is its horizontal extent. It is configured so that its width on the groundplane is identical at each row of the image. The horizontal distance that each row of the trapezoid encompasses is approximately 7.0 meters, about twice the width of a typical lane. This trapezoid is selectively sampled according to the strategy depicted in the schematic on the right of Figure 3-23 so as to create a low resolution (30x32 pixel) image in which important features such as lane markings, which converged towards to top of the original image, now appear parallel in the low resolution image. Note that this image resampling is a simple geometric transformation, and requires no explicit feature detection.

3.4.3.2.2 RALPH Processing Algorithm

The "parallelization" of road features described above is crucial for the second step of RALPH processing, curvature determination. To determine the curvature of the road ahead, RALPH utilizes an "hypothesize and test" strategy. RALPH hypothesizes a possible curvature for the road

ahead, subtracts this curvature from the parallelized low resolution image, and tests to see how well the hypothesized curvature has “straightened” the image.

3.4.3.2.1 RALPH Curvature Determination

The process RALPH utilizes to determine curvature is depicted in Figure 3-24. In this example, five curvatures are hypothesized for the original image, shown at the top. For each of the five hypothesized curvatures, the rows of the image are differentially shifted in an attempt to “undo” the curve and straighten out the image features. For left curve hypotheses, rows are shifted towards the right and for right curve hypotheses, rows are shifted towards the left. For the more extreme hypothesized curvatures (on the far left and right), the rows of the original image are shifted further than for the less extreme curvatures (in the middle). For all the hypothesized curvatures, rows near the top of the image, corresponding to regions on the groundplane further ahead of the vehicle, are shifted further horizontally than rows near the bottom of the image. This differential shifting accounts for the fact that for a given hypothesized curvature, the road will be displaced more at the top of the image, far ahead of the vehicle, than at the bottom. The exact shift distance for each row in the transformed images is determined both by the geometry of the camera and the particular curvature hypothesis being tested.




Figure 3-24: RALPH curvature hypotheses

As can be seen from Figure 3-24, the second curvature hypothesis from the right, corresponding to a shallow right turn, has resulted in a transformed image with the straightest features, and therefore should be considered the winning hypothesis. The technique used to score the “straightness” of each hypothesis is depicted in Figure 3-25. After differentially shifting the rows of the image according to a particular hypothesis, columns of the resulting transformed image are summed vertically to create a scanline intensity profile, shown in the two curves at the bottom of Figure 3-25. When the visible image features have been straightened correctly, there will be sharp discontinuities between adjacent columns in the image, as show in the right scanline intensity profile in Fig-

ure 3-25. In contrast, when the hypothesized curvature has shifted the image features too much or too little, there will be smooth transitions between adjacent columns of scanline intensity profile, as depicted in the left scanline intensity profile of Figure 3-25. By summing the maximum absolute differences between intensities of adjacent columns in the scanline intensity profile, this property can be quantified to determine the curvature hypothesis that best straightens the image features.




Figure 3-25: RALPH curvature scoring technique

An important attribute to note about this technique for determining road curvature is that it is entirely independent of the particular features present in the image. As long as there are visible features running parallel to the road, this technique will exploit them to determine road curvature. These features need not be located at any particular position relative to the road, and need not have distinct boundaries - characteristics required by systems that utilize strong a priori road models and edge detection.

3.4.3.2.2 RALPH Lateral Offset Determination

The next step in RALPH's processing is to determine the vehicle's lateral position relative to the lane center. This is accomplished using a template matching approach on the scanline intensity profile generated in the curvature estimation step. The scanline intensity profile is a one dimensional representation of the road's appearance as seen from the vehicle's current lateral position. By comparing this current appearance with the appearance of a template created when the vehicle was centered in the lane, the vehicle's current lateral offset can be estimated.

Figure 3-26 illustrates this lateral offset estimation procedure in more detail. Here, the current scanline intensity profile is depicted on the left, and the template scanline intensity profile, generated when the vehicle was centered in the lane, is depicted on the right. By iteratively shifting the current scanline intensity profile to the left and right, the system can determine the shift required to maximize the match between the two profiles (as measured by the correlation between the two

curves). The shift distance required to achieve the best match is proportional to the vehicle's current lateral offset.




Figure 3-26: RALPH lateral offset determination technique

Note that as with the curvature determination step, this process does not require any particular features be present in the image. As long as the visible features produce a distinct scanline intensity profile, the correlation based matching procedure will be able to determine the vehicle's lateral offset. In particular, even features without distinct edges, such as pavement discoloration due to tire wear or oil spots, generate identifiable scanline intensity profile variations which RALPH can exploit to determine lateral offset. This is a performance feature which edge-based road detection systems do not share.

3.4.3.2.3 RALPH Adaptation to Changing Conditions

Another important feature of RALPH stems from the simplicity of its scanline intensity profile representation of road appearance. The 32 element template scanline intensity profile vector is all that needs to be modified to allow RALPH to handle a new road type. Modifying this vector is extremely easy. In the current RALPH implementation there are four ways of adapting the template to changing conditions.

The first method involves the driver centering the vehicle in its lane, and pressing a button to indicate that RALPH should create a new template. In under 100 msec, RALPH performs the processing steps described above to create a scanline intensity profile for the current road, and then saves it as the default template. From that point on, RALPH can warn the driver of road departure danger on this road using the newly created template to determine the vehicle's position relative to the lane center.

A second method for acquiring a template appropriate for the current road type is to select one from a library of stored templates recorded previously on a variety of roads. RALPH can select the best template for the current conditions by testing several of these previously recorded templates to determine which has the highest correlation with the scanline intensity profile created for the current image.

The third method of template modification occurs after an appropriate template has been selected. During operation, RALPH slowly "evolves" the current template by adding a small percentage of

the current scanline intensity profile to the template. This allows the current template to adapt to gradual changes in the road's appearance, such as those caused by changes in the sun's angle.

RALPH handles more abrupt scene changes, such as changes in lane marker configuration, using the final and most interesting template modification strategy. In this technique, RALPH uses the appearance of the road in the foreground to determine the vehicle's current lateral offset and the curvature of the road ahead, as described above. At the same time, RALPH is constantly creating a new "rapidly adapting template" based on the appearance of the road far ahead of the vehicle (typically 70-100 meters ahead). This rapidly adapting template is created by processing the distant rows of the image in the same manner as described above. The road's curvature is assumed to be nearly constant between the foreground and background, allowing RALPH to determine where the road is ahead and therefore what the new template should look like when the vehicle is centered in its lane.

If the appearance of the road ahead changes dramatically, RALPH uses this technique to quickly create a template appropriate for the new road appearance. When the vehicle actually reaches the new road, RALPH determines that the template it was previously using is no longer appropriate, since it does not match the scanline intensity profile of the current image. It therefore swaps in the rapidly adapting template, and continues driving. Note that this rapid adaptation occurs in the time span of approximately 2 seconds, without any human intervention.

3.4.3.2.3 RALPH Performance

As part of Task 3, the project team conducted extensive laboratory, test track and on-road experiments in order to characterize RALPH's performance. The results of these tests, presented below, indicate that RALPH can accurately estimate the vehicle's lateral position on the road, as well as the curvature of the road ahead, under a wide variety of conditions.

3.4.3.2.3.1 Laboratory Tests

An important factor determining roadway departure countermeasure effectiveness is the accuracy of the sensing system employed. A system with low accuracy will be prone to false alarms, and will potentially underestimate the danger of true roadway departure situations. The crucial accuracy metric for RALPH is how well can it estimate the location of the road ahead of the vehicle, since it is the road location that will be used to determine the danger of roadway departure (See Section 3.6 for more details on this algorithm).

In order to quantify RALPH's ability to accurately determine the position of the road ahead, the project team conducted controlled laboratory tests in which accurate measurements of the road's actual location could be made. To facilitate these measurements, the team collected high quality video sequences of road scenes, using a Umatic 3/4 inch VCR. These scenes were collected in the Navlab 5 test vehicle, using the same camera mounted in the same location (next to the rear view mirror) as in the experiments described in following sections. These sequences include both day and night operation, as well as images of a variety of road types, including both rural roads and multi-lane divided highways. The test road sequences recorded on videotape were all between

four and nine miles in length. While recording the sequences, the driver repeatedly changed the vehicle's lateral position within the lane in order to obtain a wide range of images.

The video sequences were subsequently replayed in the laboratory, and RALPH was used to track the road. More specifically, RALPH combined its estimates of the vehicle's lateral offset and the curvature of the road ahead into an estimate of the lane center location one second ahead (about 25m) of the vehicle. Note this is the same estimation technique used by the ALVINN system in the previous experiments.

RALPH's lane center position estimate was compared in real time with the estimate of lane center provided manually by the experimenter. The experimenter continuously indicated his estimate of the lane center location by keeping a crosshair centered over the right lane marking one second ahead of the vehicle in the image using a computer mouse. The difference between RALPH's estimate of lane position and the experimenter's estimate was stored for later analysis.

The results of these tests are summarized in Table 3-3. For each of the conditions tested, the table shows the mean and standard deviation of the difference between RALPH's estimate of the lane center position, and the experimenter's estimated of the lane center position. In general, RALPH's performance was quite good in all the conditions tested, with the total mean disagreement between RALPH and the experimenter of 13.2cm, which is just slightly larger than the wide of a typical single lane edge marker. As was expected, lower mean and standard deviation was observed in the conditions with the most consistent features. One such situation is shown in Figure 3-27. It depicts a daytime highway scene in which the lane markers are very clearly visible. Under these conditions, the mean disagreement between RALPH and the experimenter was 11.4cm. The variance of the disagreement was 14.3cm. Note that a substantial portion of the disagreement between RALPH and the experimenter can be attributed to inconsistency in the experimenter's estimate of the lane center position. Accurately indicating the lane position 20m ahead using a mouse is a difficult task. In a series of repeatability tests, it was determined that the experimenter's estimate of lane position over two different trials on the same section of videotape varied by an average of 7.3cm.

Table 3-3: RALPH lane location estimation accuracy

Condition	Mean Error (cm)	Error Std. Dev. (cm)
Daytime Highway	11.4	14.3
Daytime Highway w/ Shadows	13.8	18.9
Nighttime Highway	11.1	13.8
Daytime Rural Road	13.7	16.2
Daytime Rural Road w/ Glare	15.8	17.2
Nighttime Rural Road	13.8	16.8
Total	13.2	16.2




Figure 3-27: RALPH processing a daytime highway image




Figure 3-28: RALPH processing a daytime highway image with heavy shadows




Figure 3-29: RALPH processing a nighttime highway image




Figure 3-30: RALPH processing a daytime rural road image




Figure 3-31: RALPH processing early morning rural road image with glare off road




Figure 3-32: RALPH processing a nighttime rural road image

On the same stretch of highway under conditions of heavy shadows (See Figure 3-27), the mean and standard deviation of RALPH's lane position estimation error increased somewhat to 13.8cm and 18.9cm, respectively. This increase in error was due primarily to the limited dynamic range of the camera, causing the shadowed regions of the image to be black and/or the areas in sunlight to be saturated.

In contrast, RALPH's lane location on the same stretch of highway ability improved slightly at night. As can be seen in Figure 3-27, the lane markers were very distinct in this situation, resulting in a mean error of 11.1cm and a standard deviation of 13.8cm.

RALPH's performance on rural roads such as the one in Figures 3-27 was fairly similar to the highway results. The mean and standard deviation under favorable daytime conditions did increase slightly over the corresponding figures for favorable daytime highway images, to 13.7cm and 16.2cm, respectively. This increase was primarily caused by two factors. First, more frequent and substantial grade changes on the rural roads changed the perspective of the camera relative to the road. This resulted in slight additional lane position estimation errors, particularly at grade transition points. Second, there were several cross streets intersecting the section of rural road tested, which occasionally resulted in momentary inaccuracy when the lane marker's disappeared. However the increase in average lane position estimation error due to these effects was small, on the order of two centimeters.

One problem with lane tracking systems which rely exclusively on lane markers to locate the road ahead is that they sometimes have difficulty when glare off the pavement makes the markers hard to find. This type of glare typically occurs when the pavement is wet, and/or when the sun is low on the horizon. To quantify the effect of these conditions on RALPH, a video sequence was collected on the same rural images during the early morning hours heading into the rising sun. An example image from this sequence is shown in Figure 3-27. As was expected, the mean and standard deviation of RALPH's error increased under these conditions, to 15.8cm and 17.2cm, respectively. However these increases were slight, again in the range of 2cm. RALPH was still able to accurately locate the road ahead under these conditions by adapting its processing to utilize the boundary between the bright pavement and the dark shoulder. This ability to adapt to changing conditions was determined to be particularly important in the on-road tests, described in Section 3.4.3.2.3.3.

In summary, the team's laboratory tests indicate that RALPH can localize the position of the road ahead of the vehicle to within approximately the width of a single lane marker under a variety of conditions. To further characterize RALPH's ability to perform repeatably and reliably, the team conducted extensive test track and on-road experiments, described below.

3.4.3.2.3.2 Test Track Experiments

Experiments were conducted at the Vehicle Research and Test Center (VRTC) in East Liberty, Ohio, and on a road segment outside of Pittsburgh often used by Carnegie Mellon for testing. These tests involved repeatedly driving the same stretch of roadway at different speeds and with different degrees of driver vigilance in order to determine whether variability in driver perfor-

mance could be detected with RALPH. These tests were performed with a single individual from the project team as the driver, when there were no other vehicles on the test road.

In the videotape experiments presented above, the goal was to quantify RALPH's ability to find the position of the road ahead by combining RALPH's estimate of the vehicle's lateral position and its estimate of the curvature of the road ahead. In the first set of test track experiments, the goal was to tease apart this combination, and measure RALPH's ability estimate the curvature of the road ahead. In this experiment, the Navlab 5 test vehicle was driven through the S-curve shown in Figure 3-33.




Figure 3-33: S-curve used for testing RALPH

Careful measurement of the first curve indicates that it has an average radius of curvature of approximately 343m. Figure 3-34 shows RALPH's estimate of the road curvature during two traversals of the entire S-curve at 55mph.



Figure 3-34: RALPH's curvature estimate on two traversals through the s-curve

Note the consistency in the curvature estimate between the two traversals. RALPH's mean estimate for radius of curvature during the first traversal of the first curve was 373m, and the mean on the second traversal was 374m. Not only are the two estimates extremely close, but they match quite closely to the measured radius of 343m. In fact, the 30m discrepancy between the measured curve radius and RALPH's estimate may at least partially be attributed to uncertainty in the manual curvature measurement.

The next set of experiments was done to determine if anomalous driver behavior can be detected using RALPH. Again the driver drove twice through the S-curve at 55mph. The first time through, the driver concentrated on accurate driving. The second time through, the driver was momentarily distracted by an in-cab task similar to the task developed for the Iowa Driving Simulator experiments (see Volume II of this report). The distractor task required the driver to glance to the back of the vehicle for up to two seconds. The goal was to determine if the lane deviations resulting from this momentary inattention could be detected in the lane tracking output RALPH produces.

A graph of RALPH's estimate of the vehicle's lateral position, both during normal driving and while the driver was performing the distractor task are shown in Figure 3-35. As can be seen from the graph, the relatively large magnitude lane deviations resulting from momentary distraction are clearly discernible when compared with driver's normal lane deviations. Algorithms to detect these anomalous lane deviations are presented in Section 3.6, and investigated further in the driver simulator tests described in Volume II of this report.




Figure 3-35: Lane deviation in normal driving, and when the driver is distracted

The results of these test track experiments indicate that RALPH can repeatably detect both the curvature of the road ahead, as well as the excessive lane deviation by the driver. However these experiments were conducted under favorable weather and lighting conditions. The next set of on-road tests were conducted in order to quantify RALPH's reliability under the range of conditions typically encountered in normal driving.

3.4.3.2.3.3 On-road Tests

One of the most significant potential drawbacks of countermeasures that rely on video cameras for sensor input is their susceptibility to adverse conditions. Systems that rely on visible features to determine the vehicle's position on the road can have trouble when these distinctions become difficult to detect, due to adverse weather, poor lighting, or degraded pavement. To quantify this effect, the project team conducted a series of on-road tests of the RALPH system.

The culmination of these experiments was a 2850 mile test drive from Washington, DC to San Diego, CA. Except for a few detours, the trip exclusively involved highway driving. The trip included many of the difficulties typically encountered in normal driving - nighttime driving, driving at sunset when the sun is low on the horizon, driving through rain storms, driving on poorly marked roads, and driving through construction areas.

During the 2850 mile trip, statistics about the RALPH system's "availability" were collected. Availability is defined to be the percent of the distance and time traveled during which RALPH was tracking the road correctly. Correct tracking is difficult to quantify in live tests (as opposed to the videotape tests described in Section 3.4.3.2.3.1). To measure tracking correctness, the assumption was made that the driver is able to steer correctly, so if the steering direction recommended by RALPH disagreed significantly from the driver's steering direction, then RALPH was not tracking correctly. In more detail, when the steering direction suggested by RALPH differed from the driver's steering direction such that following RALPH steering arc at the current speed would result in a difference in lateral acceleration of 0.04g or greater, then RALPH was judged to be tracking incorrectly. This measure is closely related to the Time-to-Trajectory-Divergence (TTD) algorithm use for lane departure warning, described in Section 3.6.

Overall, the results were quite encouraging. By the above definition, RALPH was able to accurately track the road ahead of the vehicle during 98.2 percent (2796/2850 miles) of the trip. Due to the system's ability to adapt to changing conditions, RALPH was able to track the road in situations which would be difficult for other lane trackers, particularly those that rely on finding distinct lane markers. Some of the different situations that RALPH was able to handle are illustrated in Figures 3-37 through 3-39.

Some of the roads, like the two shown in Figure 3-36, were very much like one would expect on a major highway - nice pavement and good lane markings. Even when the lane markers were missing, as on the freshly paved road in the left hand image of Figure 3-37, RALPH was able to continue tracking the road by exploiting the boundary between the pavement and the off road area. This same type of road proved quite difficult at night however, when the edge formed by the pavement boundary was no longer visible. A 10 mile stretch of new, unpainted highway encountered at night in Kansas accounted for a significant portion of the 1.8 percent tracking failure during the trip. Rain proved to be less of a problem. Even when the specular reflection off wet pavement obscured the lane markings, as in the right hand image of Figure 3-37, RALPH was able to key off the tracks left in the wet pavement by the vehicle in front to locate the road ahead.

West of the Rocky Mountains, there were some stretches of very poor roads (See Figure 3-38). Often the lane markers were nearly invisible due to wear (left). Several times there were long

stretches of construction where the road was basically very fine, packed gravel, without any lane markings (right). During these stretches, RALPH was able to exploit the differences in appearance of the packed gravel and the loose gravel around it and continue tracking the road.

The freeways in California posed an interesting challenge. Instead of having painted lane markings to delineate lanes, they had reflectors that were nearly invisible during the day (See left image, Figure 3-39). But RALPH was able to track the lane using the discoloration from the oil spot down the center of the lane. RALPH also performed well on the 1-15 HOV lane into San Diego, which had no visible lane markings, but a strong boundary between the cement road surface and the asphalt shoulder (right image, Figure 3-39).

The situation which gave the system the most difficulty was in city traffic, when the road markings were either missing or obscured by other traffic (See Figure 3-39). However, as was determined in Task 1, relatively few roadway departure crashes occur in this type of situation. Furthermore, in this and most of the other situations RALPH had difficulty with, it was able to recognize that it couldn't track the road, and inform the driver of its confusion. In a deployed countermeasure, this ability to identify confusing situations could be used to minimize false alarms.

In conclusion, extensive tests of vision-based lateral position detection systems both with and without forward preview indicate that such systems are able to accurately detect the vehicle's position and orientation relative to the roadway in a wide variety of situations. The impact on overall system performance of system inaccuracies in the few remaining conditions that do provide difficulty will be modeled as part of Task 4.




Figure 3-36: Examples of well marked roadway encountered in cross country test





Figure 3-37: Roads without strong markings (left) and with wet pavement (right)





Figure 3-38: Road with severely worn markings (left) and unpaved road (right)





Figure 3-39: California freeways with reflectors instead of painted lane markings





Figure 3-40: Challenging images from city driving

3.5 Goal 3: Inferring Driver's Intentions

Inferring the driver's intention is an important goal for minimizing false alarms in a lateral roadway departure countermeasure. For instance, what appears to be a dangerous lane excursion may actually be a harmless lane change maneuver. Other situations which could potentially result in false alarms by a lane departure warning system include:

- Pulling to the side of the road
- Avoiding an obstacle in the travel lane
- Turning onto a cross street
- Taking an exit ramp

Lane change maneuver's and pulling to the side of the road are perhaps the most difficult situations to identify as harmless, since they most resemble true roadway departure crash conditions. While no existing technology for achieving this type of discrimination could be identified for testing during our Task 3 experiments, the project team believes that such technology is implementable. The simplest approach to detecting intentional lane changes would be to monitor the vehicle's turn indicator. The obvious potential problem is that driver's often do not signal their intentions using their turn indicator. One hypothesis is that false alarms might be acceptable under these conditions, and might even encourage drivers to more faithfully utilize their turn signal.

More sophisticated methods to detect intentional maneuvers to change lanes or pull to the side of the road include monitoring for characteristic vehicle dynamic state changes. For example, lane change maneuvers typically exhibit a sinusoidal lateral acceleration pattern. Pulling to the side of the road is usually preceded by a reduction in vehicle speed. A countermeasure that could identify these characteristic signs of an intentional maneuver, could suppress warnings under these conditions and avoid false alarms.

The signs of intentional maneuvers are likely to vary from driver to driver, and therefore some form of adaptation will most likely be required to effectively avoid false alarms while ensuring that dangerous lane excursions are still detected. Promising preliminary work in this area has been conducted by Honeywell using an artificial neural network approach to discriminate between normal and dangerous driving behavior [12]. The project team recommends further investigations in this area as part of Phase II and III.

Lane excursions in order to avoid obstacles in the travel lane are an interesting case. Emergency evasive maneuvers to avoid an obstacle are relatively easy to detect, because the steering and pedal inputs provided by the driver are typically far in excess of what is usually observed during normal driving. In fact, the forward looking lateral countermeasures developed and tested in Task 3 includes a mechanism by which responses from the countermeasure are suppressed during extreme steering maneuvers. The assumption is that extreme maneuvers indicate an attentive driver, and that the judgement of an attentive driver will be more appropriate than that of a countermeasure. Further tests need to be conducted to determine if intentional extreme maneuvers can be discriminated from unintentional control inputs, such as inadvertent steering as the driver

slumps on the wheel after passing out.

Detection of intentional maneuvers to take an off ramp or turn onto a cross street should also be facilitated by monitoring the turn indicator, and by reasoning about the dynamic state of the vehicle. An additional source of information which should improve detection of these situations is a vehicle positioning system. Knowing the vehicle's location on a digital map may allow a countermeasure to infer that the vehicle is slowing down and moving towards the road edge in order to turn at the upcoming intersection. A more detailed discussion of the capabilities of vehicle position detection systems is presented in Section 4.0 on longitudinal countermeasures.

3.6 Goal 4: Detect Potential for Roadway Departure

The next action that must be performed by an effective roadway departure countermeasure is to combine the information about the vehicle and driver's state into a measure of the roadway departure danger. Two algorithms to accomplish this functional goal were investigated as part of Task 3. The first algorithm, Time-to-Line-Crossing (TLC), originally developed by Godthelp [10] uses the time until one vehicle tire will cross the lane boundary as a measure of roadway departure danger. The second algorithm, Time-to-Trajectory-Divergence (TTD), was developed as part of this program to overcome the rigidity of the TLC algorithm. TTD compares the driver's steering arc with the steering arc suggested by the countermeasure in order to determine the danger of a road departure. The results of in-vehicle tests with these two algorithms are discussed in this section. The results of driving simulator tests of these two algorithms are presented in Volume II of this report.

3.6.1 Time-to-Line-Crossing (TLC) Algorithm

The Time-to-Line-Crossing (TLC) algorithm computes the time (in seconds) until one of the vehicle's tires will cross one of the lane boundaries, if it continues along its current trajectory. The equation used to calculate TLC is extremely simple:

$$TLC = \frac{D}{V}$$

where:

D = Distance between the closest tire and the lane boundary the vehicle is moving towards (m)

V = Lateral velocity of the vehicle (m/s)

If TLC falls below a certain threshold, meaning the vehicle will shortly cross the lane boundary if it continues along its current path, a countermeasure would trigger a response to alert the driver of the danger.

The major advantage of this algorithm is that it only requires relatively easy to compute state vari-

ables, the vehicle's lateral position on the road and its lateral velocity. This makes it an appropriate algorithm for systems without forward preview, like the AURORA system described in Section 3.4.2. The team performed a number of tests of AURORA to quantify its ability to accurately calculate TLC. As was shown in Section 3.4.2, AURORA can accurately calculate one of the two parameters required for computing TLC, the vehicle's lateral position. In the tests described in this section, the vehicle's lateral velocity was calculated by AURORA using the rate of change in lateral position over time. More specifically, AURORA computed the difference between two estimates of the vehicle's lateral position, separated by a short time, typically 0.25 seconds. By dividing the change in lateral position by the elapsed time between the estimates, AURORA was able to calculate the vehicle's lateral velocity.

Using its estimates of lateral position and lateral velocity, AURORA was able to use the previous equation to estimate TLC quite accurately, as can be seen from Figure Figure 3-41. This graph was generated by comparing AURORA's estimate of the time-to-line-crossing with the actual time-to-line-crossing during the one second interval prior to crossing the lane boundary. The dark, solid line represents the actual TLC, as measured backwards from the actual time when the vehicle crossed the edge of the lane (accurate ground truth estimates of TLC were possible because the road imagery data was captured on videotape). The results of these experiments indicate that AURORA can estimate TLC with an average error of approximately 0.2 seconds, and a standard deviation of 0.23 seconds. This level of noise in the TLC estimate could potentially result in variations in warning onset of approximately 0.2 seconds. This magnitude of error should be small enough so as not to be noticed by a driver. Tests of this hypothesis were conducted in the Iowa driving simulator experiments described in Volume II of this report.




Figure 3-41: AURORA's estimate of TLC before lane crossing

The one shortcoming of the TLC algorithm that was discovered during in-vehicle tests was its tendency to produce false alarms while the driver is negotiating curves. On the approach and entrance to curves, typical drivers have a tendency to steer towards the inside of the curve. To a system without forward preview like AURORA, this can appear to be a potentially dangerous

drift towards the lane boundary. In order to avoid the false alarms that can occur in this situation, and take advantage of the additional information a system with forward preview can provide, the project team developed a second lateral warning algorithm, Time-to-Trajectory-Divergence (TTD).

3.6.2 Time-to-Trajectory-Divergence (TTD) Algorithm

Instead of computing the time until the vehicle will cross the lane boundary, the TTD algorithm compares the driver's steering arc with the "optimal" steering arc generated by the countermeasure. If the driver's and the countermeasure's arcs differ significantly, this is an indication that the driver is steering inappropriately, and that a lane departure warning should be triggered.

In more detail, the ALVINN and RALPH forward looking systems output the radius of the arc that the vehicle should follow in order to bring it to the center of the lane within a fixed time period (this period is adjustable, but 1.5 seconds is a typical value). This is presumed to be the "optimal" arc. Note that if the vehicle is off to one side of the lane, this arc represents a smooth path from the vehicle's current location back to the center of the lane. To determine whether to trigger a roadway departure warning, the TTD algorithms compares the driver's steering arc with the optimal arc. If the two arcs differ, following the driver's arc instead of the optimal arc would result in a path which diverges from the optimal trajectory. The TTD algorithm calculates the time until the two diverging paths would be a threshold distance apart (typically about 1.0m) at the current velocity. This is obviously a function of both the vehicle's speed, and the magnitude of the difference between the two steering arcs. If TTD, the time until trajectory divergence, falls below a threshold (typically about 1.2 seconds), then the countermeasure triggers a warning, since this indicates the vehicle is quickly departing the roadway. The equation for calculating TTD is given below. The geometric derivation of TTD is provided in Figure 3-42.:

$$TTD = \frac{\sqrt{2D}}{v \sqrt{\left| \frac{1}{r_c} - \frac{1}{r_p} \right|}}$$

where:

TTD = time to trajectory divergence (seconds)

r_p = turn radius of the person (1/m)

r_c = turn radius of the countermeasure (1/m)

D = threshold divergence distance (~1m)

v = vehicle velocity (m/sec)

In-vehicle tests of the TTD algorithm indicate several important characteristics. First, the TTD algorithm is able to reliably determine when the vehicle is departing the lane. TTD also appears to be less susceptible than TLC to false alarms when the driver "cuts the corner" on curves. The "optimal" trajectory as computed by the TTD algorithm naturally cuts the corner, since it involves steering the vehicle towards the center of the road a fixed distance ahead of the vehicle. Since this trajectory is similar to a person's natural trajectory, the time to trajectory divergence will remain high, even as the driver drifts towards the inside of the lane on curves.




Figure 3-42: Derivation of the Time-to-Trajectory-Divergence (TTD) algorithm

However TTD still exhibits occasional false alarms during curve negotiation, when the driver cuts the corner to a larger extent than the TTD algorithm expects. To further reduce the false alarm rate on curves, additional decision logic was added to the TTD algorithm. If the driver is cutting the curve by steering more sharply than the optimal trajectory but in the appropriate direction (i.e. left vs. right), warnings are suppressed. This allows an additional amount of drift towards the inside of curves without triggering a countermeasure response.

Although they are rare, lane departure crashes off the inside of curves account for about 20 percent of all run-off-road crashes on curves. The curve cutting logic detects departures off the inside of curves in the following manner. If the vehicle drifts too far towards the inside of the curve, the sign of the optimal trajectory changes (i.e. instead of steering left to follow the left curve, the optimal trajectory indicates the driver should steer right to return to the lane center). At this point the curve cutting logic is overridden, and the countermeasure triggers a warning. This insures that lane departures off the inside of curves will be detected by the countermeasure.

Experiments comparing the effectiveness of the TLC and TTD algorithms for assessing roadway departure danger were conducted on the Iowa driving simulator. The results of these experiments are presented in Volume II of this report

3.7 Summary

Crashes in which the vehicle drifts off the road due to inattention or driver impairment were identified in Tasks 1 to be an important category of roadway departure crashes. In Task 2, five primary functional goals were formulated for a countermeasure to prevent this type of crash. They are:

1. Monitor vehicle dynamic state
2. Determine vehicle's position/orientation relative to road
3. Infer driver's intentions
4. Detect potential for roadway departure
5. Present phased warning to driver

Experiments were conducted as part of Task 3 to assess the performance of technology for accomplishing these functional goals. While no complete lateral countermeasure which performs all of these functional goals was available for testing, the project team was able to acquire and test technology for accomplishing four of the five functional goals (all except inferring the driver's intentions).

The results of these tests were quite promising. Motion sensors for monitoring the vehicle's velocity and acceleration are inexpensive and quite reliable. Tests of vision systems both with and without forward preview indicate they are able to accurately determine the position and orientation of the vehicle relative to the roadway in all but the most extreme of conditions. Algorithms for detecting when the vehicle is in danger of departing the roadway were developed and verified. Simulator experiments described in Volume II indicate that this danger can be effectively communicated to the driver through several interfaces, include auditory and tactile signaling.

While additional work is required to further develop and quantify the performance of lateral countermeasure technology, particularly for inference the driver's intentions, it appears from these experiments that effective lane departure warning systems are possible with existing technology.