IDEA PROJECT FINAL REPORT
Contract ITS-7

IDEA Program
Transportation Research Board
National Research Council

March 31, 1995

DRIVER-ADAPTIVE WARNING SYSTEM

Prepared by:
Honeywell Technology Center
Minneapolis, MN

Project Team:
Robert Goldman, Chris Miller,
Steve Harp, and Tom Plocher

The ITS-IDEA program is jointly funded by the U.S. Department of Transportation's Federal Highway Administration, National Highway Traffic Safety Administration, and Federal Railroad Administration. For information on the IDEA Program contact Dr. K. Thirumalai, IDEA Program Manager, Transportation Research Board, 2101 Constitution Avenue N.W., Washington, DC 20418 (phone 202-334-3568, fax 202-334-3471).

INNOVATIONS DESERVING EXPLORATORY ANALYSIS (IDEA) PROGRAMS MANAGED BY THE TRANSPORTATION RESEARCH BOARD (TRB)

This investigation was completed as part of the ITS-IDEA Program, which is one of three IDEA programs managed by the Transportation Research Board (TRB) to foster innovations in surface transportation. It focuses on products and results for the development and deployment of intelligent transportation systems (ITS), in support of the U.S. Department of Transportation's national ITS program plan. The other two IDEA programs areas are TRANSIT-IDEA, which focuses on products and results for transit practice in support of the Transit Cooperative Research Program (TCRP), and NCHRP-IDEA, which focuses on products and results for highway construction, operation, and maintenance in support of the National Cooperative Highway Research Program (NCHRP). The three IDEA program areas are integrated to achieve the development and testing of nontraditional and innovative concepts, methods, and technologies, including conversion technologies from the defense, aerospace, computer, and communication sectors that are new to highway, transit, intelligent, and intermodal surface transportation systems.

The authors would like to acknowledge the contributions of Stewart Carroll and Jim Stewart at the University of Iowa Center for Computer Aided Design for their assistance in locating suitable driver simulator data for our use in this experiment and for their continued efforts to help us understand and process that data. We would also like to acknowledge the contribution of Glen Merrill, Honeywell Technology Center's Divisional Technology Manager for Honeywell's MicroSwitch Divisions, who provided the Internal Research and Development funding which financed the first phase of research under this program.

Table of Contents

1 EXECUTIVE SUMMARY

This report summarizes work on Honeywell's Driver Adaptive Warning System project funded by a TRB-IDEA grant. Our goal was to explore the application of learning algorithms to developing individualized models of a human driver's "style"—and then using these models to tailor warnings to be more appropriate and useful for the individual. We call this approach a *Driver-Adaptive Warning System* (DAWS). Such tailored warnings hold promise for reducing the number of false alarms in collision warning systems and thus enhancing the acceptance and use of such systems by drivers.

adaptive warning system approach.

Our architecture for a DAWS is presented in Figure 1. In our approach, a Maneuver Intent Recognition (MIR) module learns to discriminate different driver intents from each other on the basis of observable vehicle and world data. In operation, the MIR provides ongoing assessments of current driver intent. An Adaptive Maneuver Assessment (AMA) module references a number of stored behavioral models corresponding to different driver intent states (e.g., lane following, lane changing, traffic avoidance, etc.), which were acquired by learning from the behavior of the individual driver. Current driver behavior is compared to these models to ascertain how "normally" the driver is behaving—based on *his or her* prior actions.

The MIR and AMA modules provide their outputs, along with those a Road and Traffic Assessment module, to a Decision System which assesses whether a warning is necessary. We devised many approaches to making such decisions, but have not yet done the work to implement or fine tune them. Generally, warnings will be of two types: state-based or event-based. *State-based alarms* are provided when the driver's current *state* indicates a potentially dangerous situation-- as might be detected when fatigued or inebriated. Such situations may be detected when driver behavior fails to match the appropriate AMA behavioral models for some period of time. *Event-based alarms* are provided when a current *event* poses imminent danger to the driver. Such situations may be detected in a variety of ways including (1) when current driver intent is to perform an impossible task (e.g., change into a non-existent or occupied lane), or (2) when the AMA predicts driver behavior to be inadequate to the task at hand. Whenever the Decision System decides that a warning is needed, the Driver Interface has the responsibility to provide it.

Most of our work in this contract has focused on developing techniques to acquire and use individualized driver models appropriate to the AMA and the MIR. Using driving data collected from 38 individuals in experiments in the University of Iowa's driving simulator, we explored a variety of adaptive techniques for learning individualized models of simple lane following behavior for the AMA. Of the methods tried, the best performance was obtained using a multi-layered perceptron neural network approach with three layers of nodes operating over five inputs: lagged steering angle, lane deviation, relative heading, road curvature and vehicle velocity. Analyses performed at the end of the project showed that such networks, trained on only 5 minutes of an individual's driving data, were 25% better at predicting the future lane following behaviors for that individual than for other drivers in the sample.

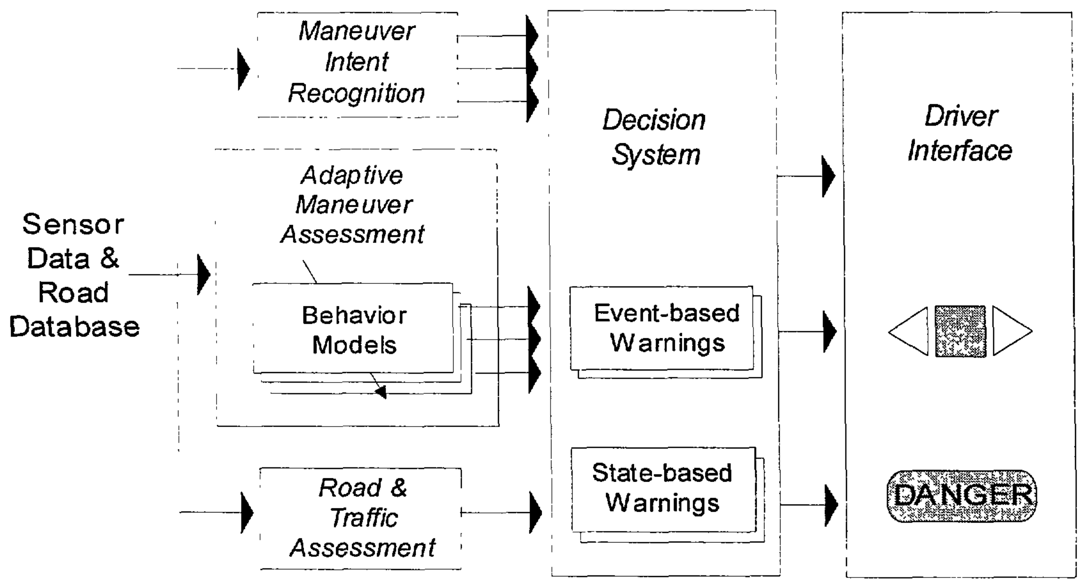

FIGURE 1 High level architecture for our driver.

Adaptive techniques were also developed for use by a MIR. We focused on the problem of discriminating lane change intents (left and right) from lane following intents in the driver data described above. Analyzing the literature and our own driver data, we developed an intent classification scheme which operates primarily on patterns occurring in relative heading angle (or yaw). Using another adaptive technique (learning vector quantization) on the relative heading data, and operating over the limited number of lane changes available, we developed a classification approach which achieved good accuracy in discriminating the three intent states. Furthermore, it appears reasonable that this discrimination can be performed in less than a second in most cases, although this depends upon allowable false alarm rates, which we have not yet studied.

Finally, we integrated our AMA and MIR modules with limited-capability versions of Road and Traffic Assessment, Decision System and Driver Interface modules into a software demonstration system which illustrates a complete, if simplified, DAWS as described in Figure 1. Using examples from our simulator data, even this simple system is capable of providing both state- and event-based alarms based on the needs of the individual driver.

The existence of large and consistent driver differences in our data serves to illustrate the need for a DAWS to augment crash countermeasures systems. Furthermore, the techniques and architectures developed on this program have paved the way to constructing various types of DAWS systems. More work is needed to validate our modeling approaches (especially with regards to detecting degraded driver states), to scale up the number of intent states-recognizable by the system, and to ensure that adequate sensor data exists to feed the DAWS. Nevertheless, this work shows that such systems are possible and hold promise for increasing driver safety.

2 PROBLEM STATEMENT

2.1 PRODUCT CONCEPT AND POTENTIAL IMPACT

The utility and effectiveness of crash countermeasure systems will depend on their acceptance and use by the driving public. If the warnings provided by these systems are perceived as unreliable or unrelated to a driver's behavior, acceptance will be minimal and the potential of the countermeasure system to reduce accidents will not be realized. Given the variability of normal driving behavior throughout the driving population, a major source of unreliable performance (i.e., false positives and false negatives) in such warning systems will be the failure of the system to correctly interpret the driver's actions. A traditional warning system would incorporate a fixed warning threshold which would be completely accurate only for the "average" or the "worst case" driver. The system will tend to misclassify the behavior of those drivers whose style deviates from the norm, either through being more cautious or more "sporty." What is needed for a successful crash countermeasure system is the ability to adapt the warning threshold to individual differences in driving style. We call a system with this capability a *Driver-Adaptive Warning System (DA WS).*

In this project, our objective was to explore machine learning techniques for acquiring an individualized model of driving behavior and to use that model to provide individualized warnings in crash countermeasure systems, thus increasing reliability and user acceptance. In our concept, a crash countermeasure system will include an on-board adaptive module that analyzes the driver's control actions (measured by on-board sensors) and, over a short period of time, develops an individualized model of the specific driver's driving style. This individually adapted model can then be used for two related purposes.

First, in conjunction with a model of the vehicle and the road surface, it can be used to make a judgment as to whether the vehicle is being operated within safe bounds, a judgment *which takes the specific driver's normal behaviors into account.* In other words, whereas a static warning system might detect a dangerous state for the "worst case" driver and sound an alarm, a DAWS would check whether this behavior was characteristic of normal, safe driving behaviors for this driver (who might be more adept) and, if so, would suppress the needless warning.

Second, the learned model will detect aberrations in a specific driver's behavior. Increased response times, more frequent or larger corrections in turning, etc. can be compared to the normal model of driving behavior for this driver to detect drowsiness, intoxication, etc. Such a capability could also provide a means of detecting when a vehicle is being driven by an unapproved driver.

For both applications, the ability to quickly, unobtrusively obtain a model of the individual's driving "style" is critical. The goal of this IDEA project was to develop and evaluate adaptive methods for acquiring this individualized model and to begin consideration of how such a model could be used in a DAWS.

2.2 INNOVATION

Our project explored the potential for a hybrid software system, based on adaptive modeling techniques, to learn patterns of individual driving behavior for use in predicting the vehicle's future state and adaptively setting alarms. Traditional warning systems (Figure 2a) accept input about the state of the vehicle, assess danger according to a rigid rule which is the same for all members of the population (e.g., lane deviation at such and such a level is dangerous), and sound an alarm whenever the rule is met. Given the variability of the driving population, warnings provided by such normative systems correlate poorly with actual behavior for many, perhaps even most, drivers. A DAWS (Figure 2b) adds individual driver sensitivity to the traditional warning system. By learning about an individual's driving behavior, we can make assessments of conditions that will (or will not) be dangerous for that individual and thus, adjust alarms so that they sound only in conditions that are appropriate for that driver. The focus of this contract is the development of an individually adaptive model for drivers. Such a component is a necessary precursor to the development of an overall DAWS, which is a logical direction for future work.

Crash countermeasure systems have focused on on-board sensors. While sensors are critical to success, rapid increases in accurate, reliable sensor technology suggest they will soon not be the main limitation on reliable performance. Rather, the greatest potential unreliability lies in the use made of those data. In future systems, drivers will be warned (or vehicle control assumed) via the system's interpretation of driver behavior and future actions. Reliable performance will depend on the accuracy of the system's interpretation of driver actions. Therefore, new approaches that take individual driver differences into account have the potential for high payoff in terms of crash countermeasure performance and acceptance.

The closest work we know on DAWS-like systems has been that of Dr. R. Onken at the Bundeswehr University in Munich (Kopf and Onken, 1992). Onken's approach is potentially more complex, harder to implement, and slower to adapt than ours. More importantly, his approach involves a single learning phase for each driver, after which the system's behavior is fixed. Our approach could provide ongoing adaptation to configure alarms not only to the driver's general behavior, but also to changes in that behavior from day to day or over a long course of driving as factors such as fatigue or boredom set in.

Lucas Advanced Engineering of Solihull and Oxford University is working with Jaguar Cars Ltd. to develop a DAWS-like system using apparently similar technology, although details are limited (SRI, 1994).

Related work is underway at Carnegie Mellon University, well known for autonomous driving research (Pomerleau, 1991; 1993; 1994). Dr. Dean Pomerleau has recently begun applying this research to warning systems (Naik, 1994) for the Department of Transportation. While Pomerleau is in an excellent position to do such research, results are apparently unavailable as yet. Shepanski and Macy (1987) developed an approach to acquiring individualized driving models similar to ours, but focused on

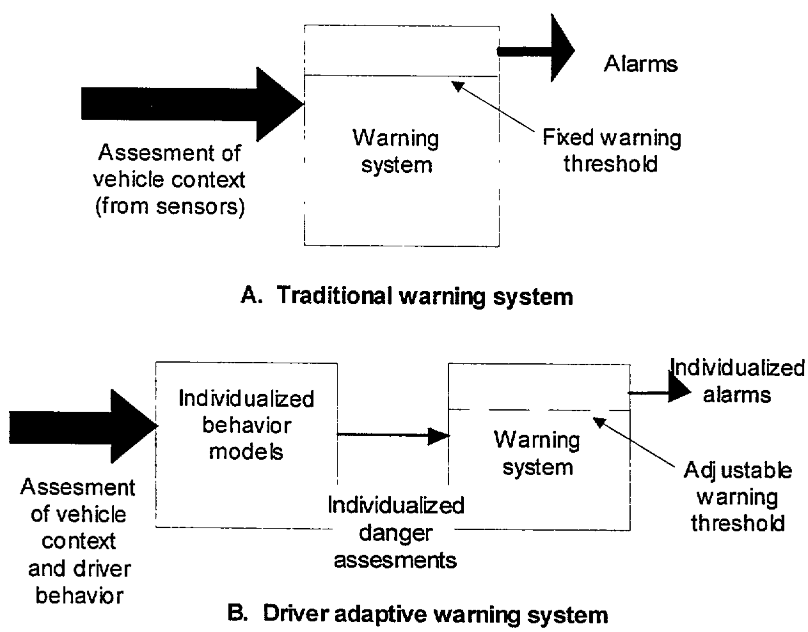

FIGURE 2 Traditional (a) and Driver-adaptive (b) Warning Systems.

autonomous control and did not use such models to provide warnings. Godthelp's (1988, 1985) work on human control of vehicles has informed our individualized modeling approach, but again, he has not addressed DAWS-like applications.

We intended to provide only an adaptive mechanism for detecting hazardous *events* (e.g., potential crashes), but our approach also shows the potential to detect persistent hazardous *states* in an individual driver (e.g., inebriation, fatigue, etc.) While there is substantial ongoing research for this latter application, our approach is innovative in being completely non-intrusive. Using only detectable vehicle parameters (e.g., wheel position, speed, lane deviation, relative heading, etc.) our approach may be able to detect deviations from a driver's normal behavior, and to discriminate the normal driving of one driver from another.

2.3 PROGRAM GOALS

The high-level architecture of our proposed DAWS was presented in Figure 1. Our goal was to learn *individualized behavioral models* for different *maneuver intent states* of a driver. Such models would populate the Adaptive Maneuver Assessment (AMA) module of our DAWS. Then we intended to develop a Maneuver Intent Recognition (MIR) model which would enable discrimination among maneuver intent states based on combinations of features from the vehicle, driver and external world. By separating the intent model from the models of driving behavior, we hoped to gain increased accuracy in these models. This would facilitate the learning process, and increase the system's flexibility in making multiple predictions given various intents. A final Decision System and Driver Interface would combine inputs from these modules to provide individualized warnings, but work on these modules was not a primary focus of this contract.

More generally, our goals for this program were to:

- Establish the existence of individual differences in driving style to evaluate the need for a DAWS,
- Demonstrate the ability of adaptive learning techniques to acquire an individualized model of driving style suitable for adapting warnings to that driver in future driving conditions-and to provide guidelines for what types of learning approaches work best for such an application.
- Establish the ability of such learning approaches to discriminate between various driver intent states (e.g., lane following versus lane changing).
- Provide an initial, high-level design for a DAWS which uses the capabilities described above.

3 RESEARCH APPROACH

3.1 PROJECT METHODOLOGY

The research plan for our IDEA project proceeded in phases to develop the proposed system architecture in Figure 1. In Phase I, supported by Honeywell Internal Research and Development funds, we began development of the individualized behavioral models of an AMA by exploring various adaptive learning approaches to acquiring behavioral models for a single, simple driver intent: lane following. We began by acquiring driving data from a variety of individuals operating the University of Iowa's driving simulator and using this data to evaluate the ability of various adaptive algorithms to predict future behavior for those drivers in similar situations.

In phase II, funded by the TRB IDEA grant, we began development of a MIR module by devising techniques for discriminating among possible intent states. We accomplished this via models capable of discriminating between lane following and lane changing intents.

3.2 TECHNICAL CHALLENGES

3.2.1 Phase I

The specific technical challenges during phase I:

1. Obtaining, interpreting and processing driving data from a variety of individuals in realistic, but appropriately constrained driving conditions.
2. Configuring learning experiments by selecting and acquiring several machine learning approaches and adapting them to work on the data collected above.
3. Interpreting these experiments to ascertain whether there were detectable, reliable differences in driving style across drivers, and ascertain which learning approaches work best to provide individualized behavioral models for an AMA from the data likely to be available to a DAWS.
4. Applying the results of the above studies via the initial, conceptual design of various complete DAWS systems and reasoning about their operations.

3.2.2 Phase II

The specific technical challenges during phase II included:

1. Improving and extending the data collected in phase I to include data for additional driver behaviors and additional data parameters as available and needed.

2. Developing an initial MIR module capable of discriminating between maneuver intents, specifically between intended lane changes and lane following.

3 Integrating this MIR capability with the AMA behavioral model for lane following developed in phase I into a unified framework for use in a DAWS.

4. Developing a software-based demonstration system to integrate our initial AMA and MIR modules into a simplified but complete DAWS system.

4 RESULTS

The primary goals of phase I were the development of an approach to acquiring individualized behavioral models for use in the AMA modules of a DAWS (cf. Figure l), and with the initial development of a DAWS architecture which would make use of this capability. Steps toward these goals are described below.

4.1 PHASE I RESULTS

4.1.1 Obtaining Data

After contract award in April, 1994, we began seeking driving data for a variety of individual drivers. Since we wanted to acquire models of individualized driving behavior in a single, simple driver intent state during phase I, we sought data for multiple drivers in straight, lane-following scenarios without the interference of other vehicles, lane changes, road curves or other hazardous road conditions.

The limited scope of this project precluded collecting new data, customized for our needs. Accordingly, we attempted to isolate appropriate data from previous experiments on the psychophysics of driving behavior. After conversations with the University of Iowa Center for Computer-Aided Design, we purchased a set of data from an experiment involving 38 subjects of varying experience driving a track involving entering and leaving a highway.

4.1.2 Organizing and Refining Data

Because this data was not customized for our needs, organizing, refining and processing it for use with our learning algorithms took a significant portion of the initial phases of this contract. One problem we encountered was the fact that the simulation laboratory at Iowa saves only the variables that are of interest for the immediate experiment; most other data is discarded.or the experiment which seemed best suited for our needs, the data which was available in the laboratory's tiles consisted of measurements of the following variables, collected at a rate of 30 Hz, for each driver: vehicle center of gravity position in three dimensions, vehicle speed, vehicle rack position, lane deviation (from center of current lane), and road type (straight or curved).

Of these, various problems in the data (described in the appendix to this report) left us with only lane deviation, velocity and rack position for use in learning models in our initial work. While still offering a feasible data set from which to predict driving behavior, note that this set does not include many variables which we might desire to make such predictions. We did not have many of the variables indicating driver intent such as use of directional indicators. We also had only measures of the vehicle's state (speed and rack position) but no direct measures of the driver's actions that affect the vehicle's state such as brake and accelerator position. We also had no indication of proximity of other vehicles, an important consideration when identifying, e.g., a passing maneuver.

Finally, the data available during this phase suffered from various types of "noise" including some uncertainty in characterizing a period of data as straight lane following (as opposed to curve following or lane changing), some basic inaccuracies in the simulator's recording of vehicle position, and some complicating factors in driver behavior such as other vehicles and roadway obstructions.

Each of these factors made the job of a learning algorithm more difficult. In order to follow this research direction further, we should collect data in experiments specifically designed to serve the needs of a DAWS: The fact that we could successfully learn individualized driving models even with this reduced data set provides an added measure of confidence for the viability of this approach in a more robust, real-world data environment.

This data organization and refinement effort took more time than expected, but the end result was the development of a data set for phase I. That set included measures of speed, rack position and lane deviation for three episodes of lane-keeping behavior for each subject.

4.1.3 Learning Approaches

The framework for learning in the state-based approach is shown in Figure 3. This is a commonly used technique in observer-based fault detection (e.g., Frank, 1991). The driver/vehicle system **is** monitored by an "observer" which compares an adaptive model to system control performance. Normally, the model tracks the behavior of the driver with only moderate errors, *(residuals)*. When the driver fails to perform as expected, residuals become greater, triggering a warning system.

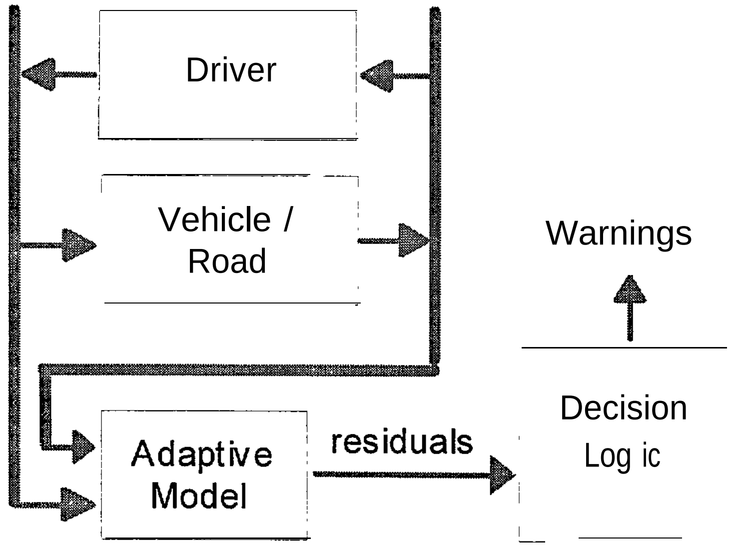

FIGURE 3 Framework for observer-based fault detection in DAWS.

During Phase I we tested several adaptive methods for acquiring individualized models of driving behavior. Specific techniques included: statistical time series techniques, multi-resolution CMAC (Cerebellar Model Articulation Controller) neural networks, and multi-layer perceptron neural networks (trained by variants of back-propagation). Results of learning experiments involving each of these techniques are presented in brief below.

In one set of trials, we used various ARMAX models (auto regressive moving average with exogenous variables) to predict rack position one second (30 time steps) into the future from observed past rack positions, velocities, and lane deviations. This learning technique can be applied in an off-line mode, as we used it, or an on-line mode with Kalman filters (Ljung and Soderstrom, 1983).

The results indicated general agreement with the driver at a coarse time scale, but substantial disagreement at finer scales. This was more evident when running the model "open loop" (without reference to actual driver responses.

Another set of trials used the cerebellar model articulation controller (CMAC) technique described by Albus (1975). This technique is often employed in real-time applications due to its rapid learning capability. The version we used employed a cascade of three CMACs at different resolutions (cf. Moody, 1989). Third order B-Splines were used as the receptive field functions. We investigated CMACs with inputs for rack position at - 1, -2 and -3 seconds, and lane deviation at -1 and -2 seconds.

The CMACs learned very quickly and were capable of providing a virtually perfect prediction of the data used to train them. Unfortunately, CMACs so trained were unable to generalize well to new examples. Various structural manipulations improved generalization somewhat, but we were not able to show generalization performance beyond the level of time series models.

A third set of trials used multi-layer perceptron networks (Hertz, Krough and Palmer, 1991) arguably the most widely used network architecture, again to predict rack position 1 second into the future. In order to capture trend information, we used as input a sliding "window" of data about the car/driver state. Each window contained three parameters: lane deviation, rack position, and velocity, for some time k.

The network was trained using conventional backpropagation via a regime where the learning rate decreased by an order of magnitude when performance reached a plateau. An example of this network's behavior on test (rather than training) data is given in Figure 7 (upper panel).

4.1.4 Analysis of Learning Approaches

Examination of driver records from the simulator indicated some interesting features of lane-following behavior. First, the data showed the existence of significant, consistent individual differences in the corrective actions taken to maintain lane position, even over the same segment of roadway. These differences are reflected in, for example, the fact that an ARMAX model trained on segment 1 for driver 1 will perform substantially better on segment 2 for driver 1 than on segment 1 for driver 2. Second, the data also showed that corrective actions are not smooth and continuous like the output of a closed-loop linear controller, but show a discrete step-like structure. This nonlinearity has been variously described. Baxter and Harrison (1979) have attempted to model it with a time lag coupled with hysteresis. While capturing some of the flavor of the response, their model proved very laborious to fit. Godthelp (1988) proposes a modal scheme. In one mode, the driver engages in error correcting closed-loop control; in the other, "error neglecting" mode, the driver operates in open loop, and allows position error to accumulate. The use of such a model requires a switching scheme to decide when to go from one mode to the other. The data presented by Godthelp suggest that several factors may influence such a switching decision..

The non-linearity of the data and complexity of the parametric models make the nonparametric techniques, such as the multi-layer perceptron network, appear relatively attractive. It is likely that real driver behavior (outside of the laboratory) will have an irreducible random component. Consequently, the precise time the driver decides to make a correction, and its precise magnitude may vary from instance to instance, even in identical circumstances. Thus, it is unlikely we will be able to formulate a perfect predictive model. Fortunately, an approximate model may still provide useful information.

4.1.5 Development of DAWS Architecture

The learned models described above are capable of representing "typical" or "normal" behaviors for a given driver. We developed DAWS architectures which use these models to (1) recognize when current behaviors for a given driver are abnormal, and (2) determine when a situation requires behaviors beyond "normal" for that driver.

Our adaptive models provide the basis for two types of alarms: *state-based* and *event-based.* A *state-based* system detects and responds to aberrant driver behavior indicative of an unsafe state (e.g., drunkenness, fatigue, inattentiveness, etc.). *Event-based* systems detect and respond to conditions or behaviors which indicate the existence of a safety-critical *event.* Designs for using our adaptive models in each type of warning system will be described below.

A state-based approach is simple using our learned models. To detect deviations from normal behavior, we would train an individualized model and use it to predict future driver/car behavior. Predictions would be matched against actual behavior to yield an error measure over time (cf. Figure 3). Significant deviations are detected when the error measure crosses a threshold. The formula for choosing an ideal threshold is the subject of future studies.

Event-based approaches work somewhat differently. A simple event-based approach would build directly on the state-based approach as illustrated in Figure 4. The large rectangle represents all possible car and driver behaviors. The outer oval represents the boundary of safe driving behaviors based on the physical capabilities of the vehicle and the road-what Onken calls the Absolute Danger Boundary (Kopf and Onken, 1992). Conditions outside this oval are inevitable crashes; conditions inside the oval might be crashes *based on driver behavior.* The goal of an event-based DAWS is to set an individualized warning threshold which provides alarms in all necessary conditions but minimizes them in unnecessary ones. This threshold would be determined on the basis of the individualized driver behavior model described above- and might be adjusted as driver state was determined to be degraded. The warning system would sound alarms whenever safety-critical *events* occurred (i.e., the driver crosses the warning threshold). Additional event-based approaches involve detecting the *need* for abnormal driver events or using the MIR to detect problematic driver intents (e.g., lane changing into a non-existent lane).

The goal of phase I was simply to develop tentative DAWS architecture designs which would make use of the learned driver behavior model. The DAWS design used for our final demonstration system (cf. Section 4.2.5 below) was drawn from among these possible architectures. Although most of these candidate designs remain conceptual, we are comfortable that the learned driver models we are developing will be useful for a variety of both state- and event-based DAWS systems.

4.2 PHASE II RESULTS

The primary phase II goal was to develop a second adaptive component to recognize driver intent-the Maneuver Intent Recognition (MIR) module (cf. Figure 1). The MIR assists a DAWS in distinguishing normal maneuvers (e.g., lane changes) from control problems.

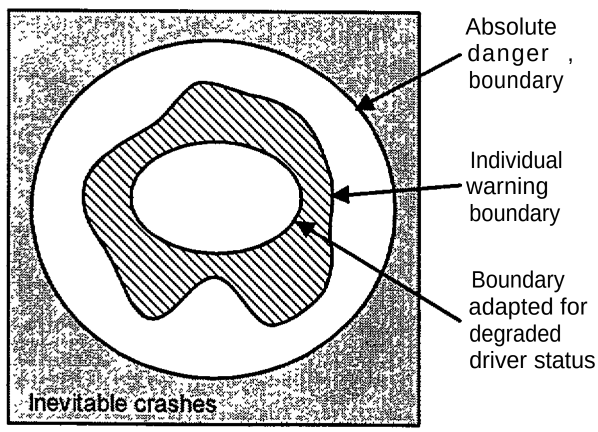

FIGURE 4 Various warning boundaries.

4.2.1 Extending Data

As a first step toward developing MIR, we worked with the University of Iowa to refme our sample data, remove noise, and provide additional data parameters such as relative vehicle heading and road curvature (Figure 5). We also obtained the relevant section of the road database used in these experiments. The course is shown in Figure 6.

Subjects entered an east-bound highway via the circular ramp in the lower left. They proceeded east and eventually exited to a northbound road. There were several suggested lane changes, but some subjects made fewer or more lane changes as suited them. Most of the driving was at highway speeds and subjects typically finished the course within about 8 to 12 minutes.

4.2.2 Improved Lane Following Assessment

With access to new data parameters, we were able to revise the learning approach from phase I. We continued to use a multi-layer perceptron approach, but now with five inputs: lagged steering angle, lane deviation, relative heading, road curvature and velocity. This approach improved the accuracy of our lane following models and extended them to cover more situations, including curves. Figure 7 (lower panel) illustrates predictions made by the phase II system (for a driver on reserved testing data) as opposed to that made by the phase I system. Prediction accuracy improved 25- 100% (depending on driver) with the new system.

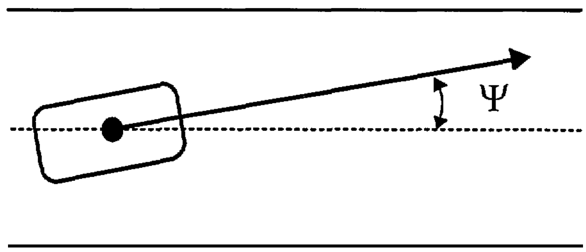

FIGURE 5 Relative vehicle heading

4.2.3 Maneuver Intent Recognition

Armed with more accurate lane following models and with lane changing data, we turned our attention to developing a Maneuver Intent Recognizer (MIR). MIR is a DAWS component that provides the decision system with context for interpreting the output of the other modules. For instance, the fit of the lane following behavioral model will tend to be poor when the driver is changing lanes. This might spark the AMA to sound an alarm, but by recognizing when the driver actually *intends* to change lanes, such a warning could be suppressed. Conversely, an intent that may cause problems, e.g., changing lanes under unsafe conditions, could trigger a warning.

There are a number of maneuvers common to normal driving but resources limited us to studying just three intent states: lane following, lane change left, and lane change right. Even for these few maneuvers there was not much simulator data available because some drivers executed only two or three lane changes.

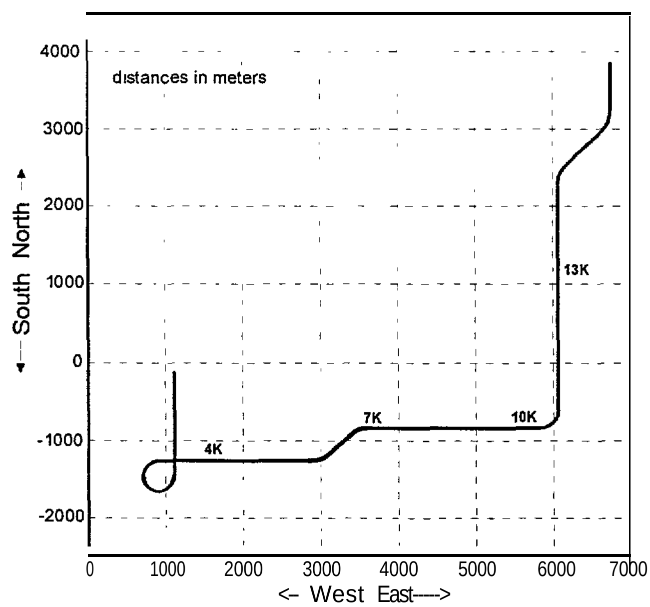

FIGURE 6 Course used in the simulator experiments. Drivers started in the lower left quadrant and proceeded to the upper right corner. The numbered hash marks indicate lane changes according to the key in the upper left.

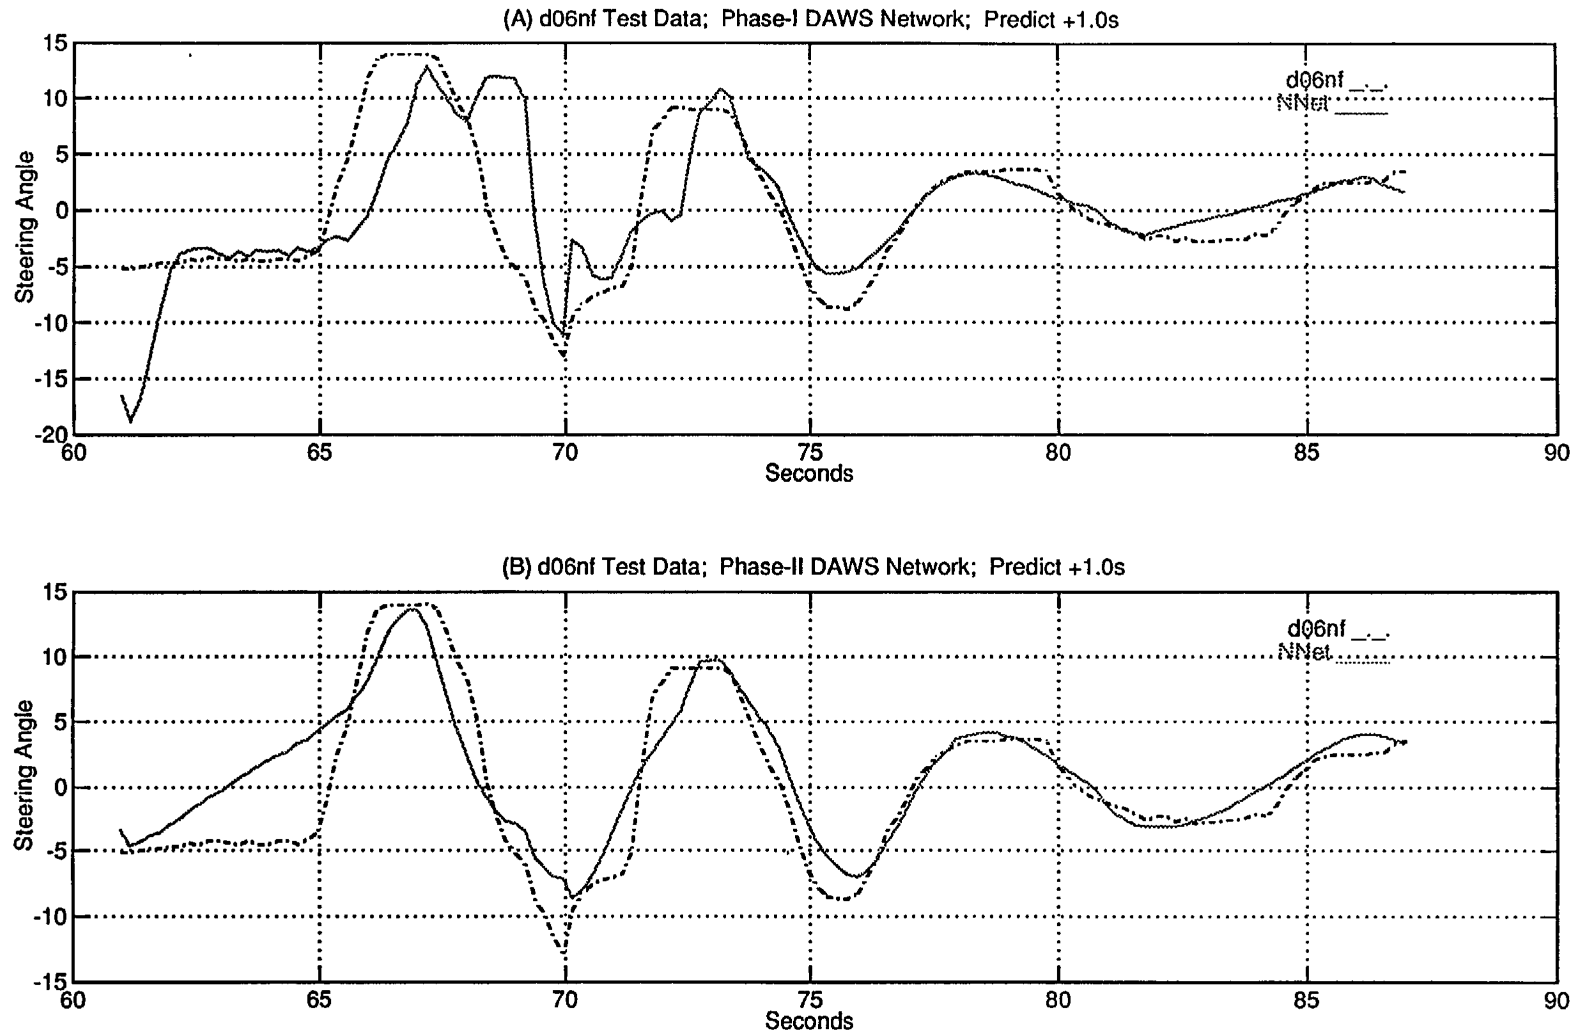

FIGURE 7 Comparison of (A) Phase-I and (B) Phase-II lane following networks on the same test sequence. Note the improved ability of the new network to anticipate driver corrections.

The current data sets do not support the individualized behavioral model development for lane following done in phase I-there are not enough instances to gain a good characterization. MIR operates by discriminating between the three intent states. When MIR decides that driver intent is lane following, the decision system compares his/her performance to the individualized lane following model to determine if the driver is behaving normally. Since there are no individualized models for lane changing, the system cannot evaluate this behavior against a learned model at this time. Ultimately, a full DAWS would see many more examples of maneuver behaviors than in our data and would learn individualized models for many common maneuvers. It would still rely on intent recognition in order to (1) decide which behavioral model to evaluate current driver behavior against, and (2) to discriminate uncommon maneuver behaviors (for which there are no models) from evaluation for state-based warnings.

4.2.3. I Qualitative Data Description

Before developing analysis methods for lane changing, we needed to understand how this behavior differs from lane following. We accomplished this through a literature review of studies of lane changing and an analysis of the instances of lane changing in our simulator data.

Simulator drivers generally showed the lane change behaviors described in the literature (e.g., Godthelp, 1985; Allen, 1982; Matshushita, et al., 1980). Qualitatively, relative heading obeyed a rough bell-shaped curve while steering angle went through a sine wave cycle. The typical maneuver was executed in about 5 seconds.

Simulator data reported which lane contained the vehicle's center of gravity at every time sample, but this proved a deceptive characterization. Lane names were changed in several places by the simulator program just as some highways get different names for different stretches. When this happened, the vehicle was assigned to a new lane without any apparent action by the driver.

Even a driver-initiated change in the lane designation was not always a true lane change maneuver. A short swerve was not considered a lane change. One driver in particular tended to ride the lane boundary, and would be credited with dozens of lane changes when in fact he merely adopted a very unusual (not to mention illegal and dangerous) lane maintenance strategy.

Although many of the lane change maneuvers followed the classic profile from the literature, others did not. A driver would sometimes execute a maneuver in two phases: an initial displacement bringing the vehicle's center of gravity near the lane boundary, and after a short delay, a second displacement to finish the maneuver. The context for such two-step operations was unpredictable, but it may relate to the presence of other vehicles. This sort of variability underlines the need to take data for these purposes from relatively natural settings; a laboratory study involving repeated executions of this maneuver under highly constrained conditions is unlikely to reveal these cases; the consequences would limit the effectiveness of any system based on the overly simplified view.

4.2.3.2 Intent Classification for Lane Changes

There are challenges in identifying lane change intent without recourse to overt indicators such as turn signals or psychological indicators like eye movements. Drivers will often adopt and maintain lane positions well off the centerline. The presence of other vehicles (data not available to us) may influence this and the early stages of a lane change maneuver look very similar to a course correction for lane following. Steering wheel angle alone is insufficient to infer intent to change lanes.

Based on our analyses, we found the single most useful predictor of lane change intent to be the relative heading angle, y (yaw). Some fairly generic predictors may be derived from examination of how it varies over time. We divided the time course into intervals at places where y becomes zero or changes sign. Within such an interval, y is then integrated (cumulative sums), so it will increase (or decrease) monotonically as the driver maintains a course heading into an adjacent lane (or off the road). The sums are reset to zero whenever y becomes zero or changes sign. This measure has a distinctive hump shape in nearly all lane changes. It is, however, nonzero at other times too. It may be refined by scaling it according to the distance to the lane in which the driver seems to be headed, thus amplifying the measure when it is executed near the intended lane boundary. Further refinements accrue by suppressing the measure when the signs of the lateral velocity and relative heading contradict the apparent lane change intent. The attack characteristics of the resulting lane change indicator also convey some information.

4.2.3.3 Using the Intent Classification Technique

To help make decisions about intent from the lane change indicator, we experimented with another adaptive technique: learning vector quantization (LVQ), (Kohonen, ***1987).*** LVQ combines a competitive layer of units that adapt to the distribution of the observed data in some input space. Each unit effectively "competes" for some portion of the input space. A subsequent linear layer is trained to make classifications based on the winning units. The input space for our intent recognition problem consisted of the indicator and its derivatives.

To train this network, we first labeled "true" lane changes in a post-session analysis. These can be found easily *after* the maneuver has occurred since they invariably correspond to large peak values of the lane change indicator. The skirts of these peaks were then taken as the maneuver boundaries. Observations in these segments were assigned categories "left" or "right" accordingly. All other observations were placed in a "following" category. The networks were then trained on subsets of these classified observations, and tested on another subset.

Our intent classification approach achieved accuracies of 92% to 98% depending on driver. Because the classification of individual observations in the first second or two of a lane change can be ambiguous, the output of the classifier can be smoothed over a number of samples (30 Hz rate) to provide greater classification confidence.

Figure 8 shows the intent recognizer output (along with lane deviation of the vehicle) for a lane change right maneuver. This case is interesting since the driver "feints left" quite strongly before making a lane change right. The feint is picked up as a partial activation of the Left indicator. Full activation of the Right indicator occurs between 423 and 428 seconds. The entire maneuver takes about 5 seconds, although the starting point is debatable.

4.2.4 Validating Adaptive Learning

Having achieved greater accuracy in our behavioral models for lane following and techniques for intent classification, we turned to some initial validation experiments for our techniques. One of the arguments for an *adaptive* warning system is that individual drivers have different styles. Furthermore, our approach to providing state-based alarms is based on the hypothesis that an individual's driving style changes as s/he becomes drowsy, inebriated, etc.

Various researchers have demonstrated individual differences between drivers (e.g., Lechner and Perrin, 1993) and differences within a drivers' behavior under abnormal conditions. For example, Wierwille and Muto (1981) noted that "Fresh drivers use small, precise corrections where necessary [but] as time passes, drivers use coarser corrections." Similarly, Mot-timer and Sturgis (1981) showed control performance of drivers under the influence of alcohol as similar to that of novices and claimed "This is further shown by the decrease in the mean path angle and mean yaw rate frequency bandwidths, which were decreased in all the alcohol dose conditions."

To show that our lane following models could detect such differences is difficult with our existing simulator data, since all the subjects remained in control of the vehicle for the entire course. They were not known to be impaired by alcohol, fatigue or unusual conditions.

However, it was still possible to indirectly evaluate the utility of the approach by demonstrating the ability to detect individual differences between drivers. To this end, we trained lane following models on half of the data for each of eight drivers and then tested them on the remaining half-either from the same or a different driver. Of interest was whether the model trained on a given driver is superior for *that* driver than models trained for others. The performance measure was RMS prediction error in steering angle. An analysis of variance tested this hypothesis in the resulting comparisons. The single df interaction for same vs. different was highly significant: $F(l,48) = 15.3$, $p<0.0005$. The lane following model was, on average, about 1.0 degree of steering angle (or about 25%) superior on the driver for which it had trained.

A similar between-drivers test was carried out for the lane change classifier model developed for our MIR. This test also showed a significant ($p < 0.005$) performance increment for test sets of drivers on which they had trained. However, the small number of lane change observations makes this conclusion highly tentative-it is unclear whether we have captured a general characteristic of the driver or something specific to this trip.

4.2.5 Demonstration System

The development of the MIR in phase II enabled us to implement a full, if highly simplified, DAWS as described in Figure 1 as a demonstration of our concept. The AMA for this software prototype contains a single individualized behavioral model, that for lane following. The MIR can discriminate between three driver intent states: lane following and lane changes left and right. The Road and Traffic Assessment module contains only the model of the simulator course as presented in Figure 6.

Each of these modules provides inputs to the Decision System. Although simplified, our Decision System is still capable of providing both state-based and some event-based warnings. The Decision System uses the intent classification provided by the MIR to decide whether or not to compare current driver behavior against the predictions of the lane following behavioral model. If current driver behavior while lane following does not match predictions based on the normal lane following model for that driver, then a state-based alarm can be sounded. The optimal criteria

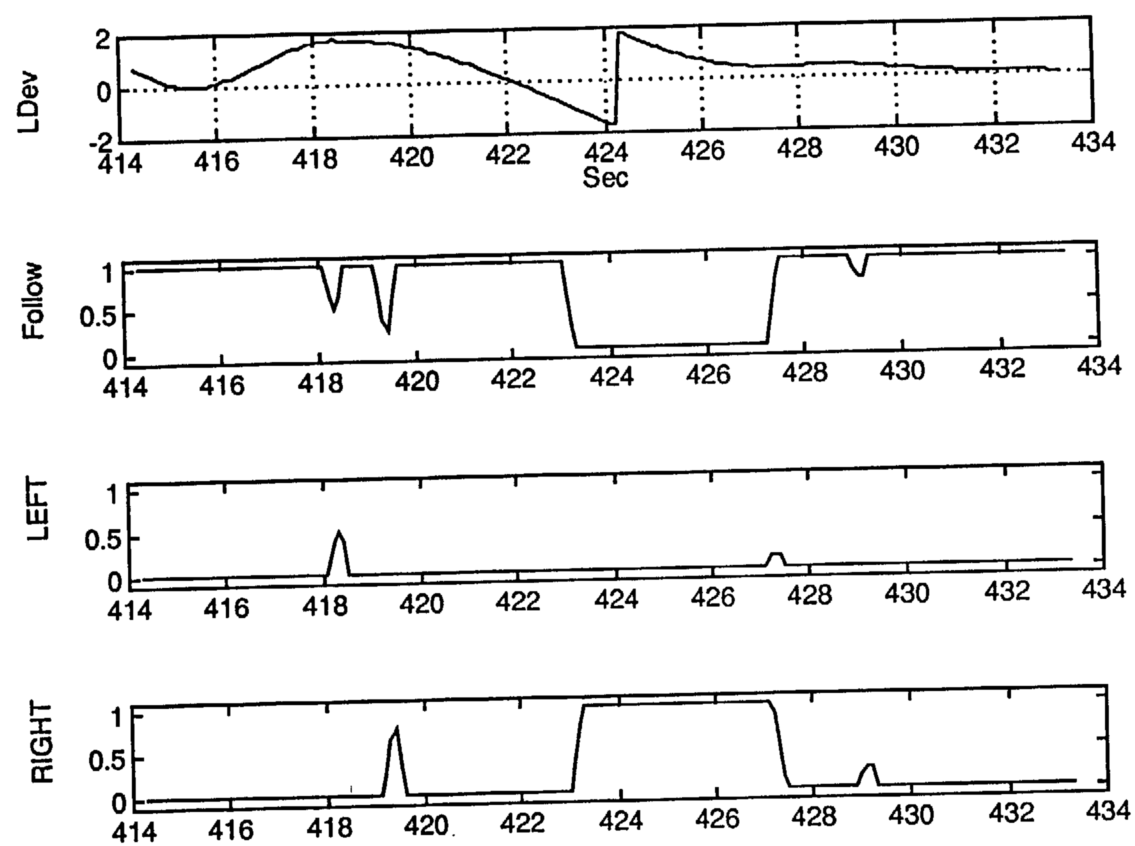

FIGURE 8 Lower three strips: Intent indicators during a lane change right maneuver following, lane change left, lane change right. Top strip is lane deviation; vehicle Cg crosses the line just past second 424.

for performing this comparison has not been finalized in this research and will require extensive human subjects testing for reliability. As noted above, since we do not have access to degraded-state driving data, the only way we can illustrate this warning approach is by using a behavioral model trained for one driver to assess the data from another driver.

A limited-form of event-based alarm can also be provided using the outputs of the MIR and the Road and Traffic Assessment modules. Whenever the MIR detects a lane change intent, the Decision System checks the road model to see if there is, in fact, a lane in the direction of the lane change. If not, it infers a run-off-road event in progress and sounds an event-based alarm to correct it.

To illustrate this prototype DAWS, we developed demonstration software that monitors vehicle position, heading and lane deviation data from a simulator trace as the driver traverses the course, and presents the output of the AMA and MIR modules. The demonstration software also provides "warning" outputs according to the decision criteria described above. It allows a user to zoom in on any course segment and display a selected subset of over a dozen variables related to vehicle state. The system is written in Matlab (by Mathworks Inc.), the mathematical visualization package used for all of our analyses in Phase II. It should be portable to any platform with the Matlab software, e.g., PC, Macintosh, HP, Sun.

5 CONCLUSIONS

5.1 GENERAL CONCLUSIONS

This project has demonstrated the existence of significant individual differences in driving style. Drivers differ widely and consistently in many parameters including: preferred speed, lane position, frequency and amplitude of steering corrections, and timing and rate of lane changing.

We also showed that various adaptive learning algorithms were able to acquire models of such differences. There are strengths and weaknesses to these approaches. We have tested a range of techniques and parameters sufficient to be confident that our selected approach provides a good mix of features for the driving domain.

Finally, we have shown that our learned models can be used in various ways to produce both state- and event-based DAWS. Further work is needed to develop such systems, but our efforts show that such systems are feasible and provide guidelines on how to construct them.

5.2 BREAKTHROUGHS AND INNOVATIONS

Major innovations in our work on this program include:

- The demonstration of adaptive modeling techniques capable of learning individualized models of driving behavior sufficiently robust to enable superior predictions of future driving behavior for the same individual. To the best of our knowledge, only Onken (1992) has attempted a similar approach and our techniques are substantially different from his.
- The development of a framework of multiple intent states and the ability to discriminate between some of them using only non-intrusive road and vehicle data.
- The integration of both individualized behavioral models and an intent classification system into a driver-adaptive warning system architecture, and the development of approaches to using this very non-intrusive approach to provide *both* state-based and event-based alarms to drivers.

5.3 APPLICATIONS AND PRODUCT RELEVANCE

The product we have been working toward is a driver-adaptive warning system which, as illustrated in Figure 1, would reside along with the vehicle-state and world state sensors and with the alarm interface on an advanced automobile, truck or other human-controlled vehicle. This DAWS would read vehicle sensors to learn models of the different authorized drivers of that vehicle and would then use this knowledge to adjust its alarm policy to best suit the needs of the individual drivers. This technology is expected to make all types of driver-dependent alarms (both event-based: run-off-road, crash avoidance, overspeed, vehicle following, etc. and state-based: fatigue, inebriation, inattention, etc.) more focused and pertinent to the individual driver. This should, in turn, result in greater driver acceptance and greater safety.

Many significant issues remain to be addressed before such a system can be productized. The highest priority issue is that a DAWS must have access to substantial vehicle and road sensor data. These data are not currently available from either the vehicle or the highway. While the necessary sensors can be added to the vehicle by original equipment manufacturers or as after-market products, the roadway markers or references that they sense will probably require changes to the highway infrastructure. These will likely be in the form of magnetic tape or reflective markers laid down on or into the pavement. Many advanced highway concepts are under investigation that share this need for a road reference. Still, in order for a DAWS to be realized, the public sector will have to make a commitment to provide the necessary infrastructure. A second and critical issue of a public nature is that of the potential liability that will be incurred by the manufacturer of a DAWS-like system. In the current legal environment of the United States, this serves as a strong deterrent to companies that might otherwise more aggressively pursue DAWS and related technology development.

The remaining issues in the productization of a DAWS are largely technical. With the appropriate resources applied to the problem, we believe that they all are resolvable over approximately a five year development time frame. These technical issues are discussed in detail below in Section 5.4.

In summary, Honeywell has no plans to market a DAWS at this time, though the techniques developed here are being considered for application to our aircraft warning systems, to individualizing operator support systems in industrial processing domains, and for their ability to provide individual human performance models across our extensive controls research and development activities. In addition to these internal activities, Ford Motor Company has expressed interest in this program and we will continue to work with them, as well as others, to find ways to bring this work closer to market and application.

5.4 FUTURE WORK TOWARD DAWS IMPLEMENTATION

The goal of this program was to develop approaches to learning individualized driver behavior models for a DAWS. We have successfully accomplished this goal. Future work should extend and refine this modeling approach, while also addressing the system issues--particularly in the sensor area-- that need to be resolved for a DAWS to become a reality. The steps envisioned for DAWS implementation are outlined below.

The first step in furthering the evolution of this technology will be to validate, on a larger, more varied driver database, the learning approaches we have already developed. For example, we know as a result of this study that our current algorithms, using only a few vehicle parameters over five minutes of driving can acquire or learn a model that significantly enhances our ability to predict driver behavior. Exactly how good these predictions can get, given more extended periods of learning and more parameters, remains a key question to be resolved. Also, we must extend our approach to additional and more varied intent states. A taxonomy of driver intent states will have to be developed and a behavioral model acquired for each of them. The system's task of discriminating among intent states becomes more difficult as the number of possible states increases, so the development of this taxonomy and the models to match it will be a critical task for realizing a DAWS.

A fundamental prerequisite to the further modeling developments described above is the availability of more extensive driver data. It will be necessary to collect data from studies specifically designed for our purposes. Especially important will be the collection of more complete driver action data (steering wheel and accelerator inputs, lane change signal actions, etc.) and world state data (locations and attributes of neighboring vehicles on the highway, etc.). Also, the collection of data for individual drivers over an extended period of time will be essential to assess the relationships between model acquisition or adaptation time and predictive power. Finally, we will need data showing drivers who are not at their best (fatigued, drunk, etc.), preferably the same driver under both conditions. Examples of naturalistic run-off-road incidents also will be desirable.

Where will the data come from? Experiments conducted in a driving simulator appear to be the best option during this next stage of development. Simulators provide the opportunity to easily record all the parameters we need for future developments. By their nature, they also can provide "ground truth" information describing the position of vehicles on the roadway, not to mention the precise location of the roadway itself. Simulators also have the advantage of allowing experiments with impaired driver states to be conducted safely. Run-off-the-road incidents can be observed with no harm to people or property.

The only question here is whether or not a driving simulator can provide sufficient fidelity in the dynamics and control responses of the vehicle to generate data that is sufficiently realistic. In order to mitigate this concern, we would recommend continued use of the 6-DOF motion base Iowa Driving Simulator to support future DAWS developments. This simulator is based on the Real-Time Recursive Dynamics (RTRD) modeling environment, recently developed at Iowa under NHTSA funding, to provide extremely high fidelity dynamics in vehicle simulations, Although validation of the Iowa Driving Simulator dynamics has not been completed, it appears, from our experience, to generate data of a form and level of complexity very comparable to real driving.

Refining the necessary sensor technology to support a DAWS will be another major step in the development and implementation process. These sensor developments can be done in parallel with the model development activities described above. Critical to the DAWS models will be a reliable way to sense road references (either center or edge). From these, DAWS will extract the relative heading measure that we have found to be fundamental to our predictions. Magnetic sensing technology holds great promise as a cost-effective, all-weather approach to this problem. Further development of this technology in order to obtain the accuracy required by a DAWS should be a high priority.

Other vehicles in the vicinity of the subject vehicle, their speed and acceleration, also should be sensed for DAWS to recognize maneuver intents and determine safety. Although the technology for these proximity-type sensors already exists, the specific DAWS requirements for them will have to be defined.

Finally, having completed the necessary mode! refinements and developed the necessary sensor technology, an extended series of driver performance and acceptance tests will have to be conducted. The initial test of the system should probably take place within the safe and controlled confines of the driving simulator, where situational and driver variables can be readily manipulated to challenge the system. Later tests must involve the complete DAWS, interacting with real sensors on a real vehicle in real traffic.

6 REFERENCES

Allen, R. (1982) Stability and performance analysis of automobile driver steering control. SAE paper 820303, Society of Automotive Engineers, Detroit, USA.

Albus, J.S. (1975) A new approach to manipulator control: the cerebellar model articulation controller (CMAC). *J. Dyn. Sys. Meas., Contr., 97 (220).*

Baxter, J., Harrison, J. Y. (1979) A nonlinear model describing driver behavior on straight roads. Human Factors, 2 1, 1, 87-97.

Frank, P.M. (199 1) Enhancement of robustness in observer-based fault detection. *Proc. IFAC/IMACSA Symposium*, Baden-Baden, Sept. 199 1. R. Isermann, B Freyermuth Editors. Oxford: Pergammon.

Godthelp, H. (1985) Precognitive Control: open- and closed-loop steering in a lane-change manoeuver. *Ergonomics,* 28, 10, 1419-1438.

Godthelp, H. (1988) The limits of path error-neglecting in straight line driving. *Ergonomics, 31,4, 609-620.*

Hertz, J., Krogh, A., Palmer, R.G. (199 1) *Introduction to the Theory of Neural Computation.* NY: Addison Wesley.

Kohonen, T. (1987) *Self-Organization and Associative Memory*, 2nd Edition. Berlin: Springer.

Kopf, N, Onken, R. (1992) DAISY, A knowledgeable monitoring and warning aid for the driver on Germany Motorways, *Proceedings of the IFAC Man-Machine Symposium.*

Lechner, D., Perrin, C. (1993) The actual use of the dynamic performance of vehicles. *Journal of Automobile Engineering, 207, 249-256.*

Ljung, L., Soderstrom, T. *(1986) Theory and Practice of Recursive Identification.,* 3rd Edition. Cambridge, MA: MIT Press. Matsushita, A., Takanami, K., Takeda, N., Takahashi, M. (1980) Subjective evaluation and vehicle behavior in lane-change maneuvers. SAE Technical Paper 800845.

Moody, J. (1989) Fast learning in multi-resolution hierarchies. In *Advances in Neural Information Processing Systems* I, David Touretzky Ed., NY: Morgan Kaufmann.

Mortimer, R.G., Sturgis, S.P. (1981) Effect of low and moderate levels of alcohol on steering performance. In *Proceedings of the First European Annual Conference on Human Decision Making and Manual Control.* H.G. Stassen, Editor. pp. 329-344.

Naik, G. (1994) This robot drives and may save lives. *Wall Street Journal,* April 4, B l-2.

Pomerleau, D.A. (199 1) Efficient training of artificial neural networks for autonomous navigation. *Neural Computation, 3, 88-97.*

Pomerleau, D.A. (1993) Neural networks for intelligent vehicles. *Proceedings of the Intelligent Vehicles '93 Symposium,* Tokyo, 19-24.

Pomerleau, D.A. (1994) Reliability estimation for neural network based autonomous driving. *Robotics and Autonomous Systems* 12, 113-l 19.

Shepanski, J., Macy, S. (1987) Manual training techniques of autonomous systems based on artificial neural networks. *Proceedings of the First IEEE Conference on Neural* Networks, *4, 697-704.*

SRI International (1994) Tech monitoring, Neural Networks, Dec. 1994/Jan. 1995,3.

Wierville, W.W., Muto, W.H. (198 1) Significant changes in driver-vehicle response measures for extended duration simulated driving tasks. In *Proceedings of the First European Annual Conference on Human Decision Making and Manual Control.* H.G. Stassen, Editor. pp 298-3 14.

APPENDIX: DRIVING SIMULATOR DATA ISSUES

The driving simulator data purchased from the University of Iowa proved to be both more and less than expected. Organizing, refining and processing this data for use with our learning algorithms took a significant portion of the initial phases of this contract. Hence, the nature of the data and the problems we encountered with it will be discussed in some detail below.

Our first task in working with the Iowa data was to isolate relevant sections from each driver's experimental sessions which illustrated simple lane-following behavior on straight road segments. One problem we encountered early on was the fact that the simulation laboratory at Iowa does not save all of the data collected from any given simulation run. In fact, only the variables that are of interest for the immediate experiment are saved; most other data is discarded. For the experiment which seemed best suited for our needs, the data which were available in the laboratory's files consisted of measurements of the following variables, collected at a rate of 30 Hz, for each driver:

- vehicle center of gravity position in three dimensions
- vehicle speed
- vehicle rack position
- lane deviation: vehicle deviation from center of current lane
- road type: straight or curved

Of these, the center of gravity position was of little use, because we did not have the position of the center of the current lane. Accordingly, there was no way to relate position in three-dimensional space to lane deviation. As a result, we were unable to make use of raw position in our experiments (although the raw position has helped us to identify sections of the data representing lane-following behavior). The road type variable was not computed reliably so that variable also could not be used.

The variables we were left with, then, were lane deviation, velocity and rack position. While still offering a feasible data set from which to predict driving behavior, note that this set does not include many variables which we might expect to improve those predictions. We do not have many of the variables that would indicate driver intent such as use of directional indicators. We also have only measures of the vehicle's state (speed and rack position) but no direct measures of the driver's actions that affect the vehicle's state such as brake and accelerator position. Finally, we had no indication of proximity of other vehicles on the road, obviously an important consideration when attempting to identify, e.g., a passing maneuver.

Each of these factors had the effect of making the job of a learning algorithm more difficult. In order to follow this research direction further, we should collect data in experiments specifically designed to serve the needs of the warning system project. More details on this topic will be provided in section 3 below.

The experiments reported here required data that described drivers trying to stay in a lane on straight roads; but the data we obtained showed this behavior mixed with lane changes, curves, etc. Since all of the subjects followed one of two tracks, we were able to use the center of gravity information to determine sections of the track which represented relatively straight roadway. For most subjects, there were three such sections. We have extracted data for those sections for each subject. This is essentially the data over which we have trained and evaluated the various learning methods described in the main body of this report.

Two additional problems needed to be resolved. First, these segments still contained some undesirable driving behaviors, in particular, some lane change behaviors. We developed a simple heuristic filter to identify such maneuvers, and then removed them from the segments.

A second problem was the presence of substantial discontinuities and noise in the data. There was a large and systematic noise component in the lane deviation data owing to choices made in the simulation software which determined the simulated car's position relative to the roadway. This was manifested as instantaneous discontinuities and transpositions in the lane deviation data.

We developed a simple model of the kinds of errors that were present and then applied an interpolating filter on the data to remove the discontinuities. The process of identifying the noise (and, in particular, distinguishing noise from valid lane changes) was quite time-consuming, as was the development of the filter. The application of the filter was also labor-intensive; since we were filtering a derived measure, the filter was not uniformly successful.

This data organization and refmement effort, which we had anticipated would take roughly one month, instead consumed four months of the project. The end result of this effort was the development of a data set for the first phase experiment. That data set included measures of speed, rack position and lane deviation for three episodes of lane-keeping behavior for each of thirty-eight subjects.