U.S. Department of Transportation

Federal Highway Administration

Publication No. FHWA-RD-94-092
July 1995

Description of the Integrated Driver Model

Research and Development
Turner-Fairbank Highway Research Center
6300 Georgetown Pike
McLean, Virginia 22101-2296

FOREWORD

This report describes a simulation model for predicting driver behavior and system performance when the automobile driver performs concurrent steering and auxiliary in-vehicle tasks. This model is an integration of two previously existing computerized models referred to as the "procedural model" and the "driver/vehicle model." The procedural component deals primarily with in-vehicle tasks and with the task-selection and attention-allocation procedures, whereas the driver/vehicle component predicts closed-loop continuous control (steering) behavior. Given descriptions of the driving environment and of driver information-processing limitations, the resulting integrated model allows one to predict a variety of performance measures for typical scenarios. These measures include time histories for vehicle state variables, such as lane position and steering wheel deflection, as well as for allocation of visual and cognitive attention. Model calibration and validation are discussed, and use of the model in analyzing complex task situations and in generating human factors guidelines is demonstrated.

Lyle Saxton, Director
Office of Safety and Traffic
Operations Research and Development

NOTICE

Technical Report Documentation Page

1. Report No. FHWA-RD-94-092	2. Government Accession No. PB95-241568	3. Recipient's Catalog No.
4. Title and Subtitle DESCRIPTION OF THE INTEGRATED DRIVER MODEL		5. Report Date July 1995 6. Performing Organization Code
7. Author(s) William H. Levison and Nichael L. Cramer		8. Performing Organization Report No. BBN Report No. 7840
9. Performing Organization Name and Address BBN Systems and Technologies 10 Moulton Street Cambridge, Massachusetts 02138		10. Work Unit No. (TRAIS) 3B3B2032 11. Contract or Grant No. DTFH61-89-C-00044
12. Sponsoring Agency Name and Address Office of Safety and Traffic Operations R&D Federal Highway Administration 6300 Georgetown Pike, McLean, Virginia 22101-2296		13. Type of Report and Period Covered Final Report Dec. 1989 - Nov. 1993 14. Sponsoring Agency Code

15. Supplementary Notes
Contracting Officer's Technical Representative (COTR): Nazemeh Sobhi, HSR-30

This research was jointly funded by the Federal Highway Administration and the National Highway Traffic Safety Administration.

16. Abstract
A simulation model for predicting driver behavior and system performance when the automobile driver performs concurrent steering and auxiliary in-vehicle tasks is described. This model is an integration of two previously existing computerized models referred to as the "procedural model" and the "driver/vehicle model." The procedural component deals primarily with in-vehicle tasks and with the task-selection and attention-allocation procedures, whereas the driver/vehicle component predicts closed-loop continuous control (steering) behavior. Given descriptions of the driving environment and of driver information-processing limitations, the resulting integrated model allows one to predict a variety of performance measures for typical scenarios. These measures include time histories for vehicle state variables such as lane position and steering wheel deflection as well as allocation of visual and cognitive attention. Model calibration and validation are discussed, and use of the model in analyzing complex task situations and in generating human factors guidelines is demonstrated.

17. Key Words IVHS, driver model, driver information system, human factors, engineering, attention.	18. Distribution Statement No restrictions. This document is available to the public through the National Technical Information Service, Springfield, VA 22161.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	124	

Form DOT F 1700.7 (8-72)
(PF V2.1, 12/13/93)

Reproduction of completed page authorized

PREFACE

The United States Department of Transportation (DOT), through its Intelligent Transportation Systems (ITS) program, is developing solutions to the most pressing problems of highway travel. The goal is to improve traffic operations and reduce congestion, accidents, and air pollution from vehicles, by applying computer and communications technology to highway transportation.

If these systems are to succeed in solving the Nation's transportation problems, they must be safe and easy to use, with features that enhance the driving experience. The University of Michigan Transportation Research Institute (UMTRI), under contract to DOT, has undertaken a project to help develop driver information systems for cars of the future—one aspect of ITS. This project concerns the driver interface—the controls and displays that the driver interacts with, as well as their presentation logic and sequencing. This is 1 of 16 reports that document that work.

This project had three objectives:

- Provide human factors guidelines for the design of in-vehicle information systems.
- Provide methods for testing the safety and ease of use of those systems.
- Develop a model that predicts driver performance in using these systems.

Although only passenger cars were considered in the study, the results apply to light trucks, minivans, and vans. Another significant constraint was that only able-bodied drivers were considered. Disabled drivers are likely to be the focus of future DOT research. A complete list of the project reports and other publications is included in the final overview report (Paul Green, 1993, *Human Factors of In-Vehicle Driver Information Systems: An Executive Summary*, Technical Report No. UMTRI-93-18, Ann Arbor, MI: The University of Michigan Transportation Research Institute). A brief summary of the work appears in Green, Williams, Serafin, and Paelke (1991).

To put this report in perspective, the project began with a literature review and focus groups to examine driver reactions to advanced instruction. Subsequently, the extent to which various driver information systems might reduce accidents, improve traffic operations, and satisfy driver needs and wants, was analyzed. That analysis led to the selection of two systems for detailed examination (traffic information and cellular phones) and contractual requirements stipulated three others (navigation, road hazard warning, and vehicle monitoring). Each system was examined separately in a sequence of experiments. In a typical sequence, patrons at a local driver licensing office were shown mockups of interfaces to determine driver understanding of the interfaces and preferences. Interface alternatives were then compared in laboratory experiments involving response time, driving simulation, and other methods. The results for each system are described in a separate report. To check the validity of those results, several on-road experiments were conducted to obtain performance and preference data for the various interface designs.

In parallel with this work, UMTRI developed test methods and evaluation protocols, UMTRI and Bolt Beranek and Newman Inc. (BBN) developed design guidelines, and BBN worked on development and application of the driver model. The driver model is the subject of this report.

SI* (MODERN METRIC) CONVERSION FACTORS

APPROXIMATE CONVERSIONS TO SI UNITS

Symbol	When You Know	Multiply By	To Find	Symbol
		LENGTH		
in	inches	25.4	millimeters	mm
ft	feet	0.305	meters	m
yd	yards	0.914	meters	m
mi	miles	1.61	kilometers	km
		AREA		
in^2	square inches	645.2	square millimeters	mm^2
ft^2	square feet	0.093	square meters	m^2
yd^2	square yards	0.836	square meters	m^2
ac	acres	0.405	hectares	ha
mi^2	square miles	2.59	square kilometers	km^2
		VOLUME		
fl oz	fluid ounces	29.57	milliliters	mL
gal	gallons	3.785	liters	L
ft^3	cubic feet	0.028	cubic meters	m^3
yd^3	cubic yards	0.765	cubic meters	m^3
NOTE: Volumes greater than 1000 l shall be shown in m^3.				
		MASS		
oz	ounces	28.35	grams	g
lb	pounds	0.454	kilograms	kg
T	short tons (2000 lb)	0.907	megagrams (or "metric ton")	Mg (or "t")
		TEMPERATURE (exact)		
°F	Fahrenheit temperature	5(F-32)/9 or (F-32)/1.8	Celcius temperature	°C
		ILLUMINATION		
fc	foot-candles	10.76	lux	lx
fl	foot-Lamberts	3.426	candela/m^2	cd/m^2
		FORCE and PRESSURE or STRESS		
lbf	poundforce	4.45	newtons	N
lbf/in^2	poundforce per square inch	6.89	kilopascals	kPa

APPROXIMATE CONVERSIONS FROM SI UNITS

Symbol	When You Know	Multiply By	To Find	Symbol
		LENGTH		
mm	millimeters	0.039	inches	in
m	meters	3.28	feet	ft
m	meters	1.09	yards	yd
km	kilometers	0.621	miles	mi
		AREA		
mm^2	square millimeters	0.0016	square inches	in^2
m^2	square meters	10.764	square feet	ft^2
m^2	square meters	1.195	square yards	yd^2
ha	hectares	2.47	acres	ac
km^2	square kilometers	0.386	square miles	mi^2
		VOLUME		
mL	milliliters	0.034	fluid ounces	fl oz
L	liters	0.264	gallons	gal
m^3	cubic meters	35.71	cubic feet	ft^3
m^3	cubic meters	1.307	cubic yards	yd^3
		MASS		
g	grams	0.035	ounces	oz
kg	kilograms	2.202	pounds	lb
Mg (or "t")	megagrams (or "metric ton")	1.103	short tons (2000 lb)	T
		TEMPERATURE (exact)		
°C	Celcius temperature	1.8C + 32	Fahrenheit temperature	°F
		ILLUMINATION		
lx	lux	0.0929	foot-candles	fc
cd/m^2	candela/m^2	0.2919	foot-Lamberts	fl
		FORCE and PRESSURE or STRESS		
N	newtons	0.225	poundforce	lbf
kPa	kilopascals	0.145	poundforce per square inch	lbf/in^2

* SI is the symbol for the International System of Units. Appropriate rounding should be made to comply with Section 4 of ASTM E380.

(Revised September 1993)

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION AND OVERVIEW

INTRODUCTION

This report describes the integrated driver model for predicting automobile driving performance that was developed in support of the experimental study, performed by the University of Michigan Transportation Research Institute (UMTRI) for the U.S. Department of Transportation, to develop evaluation methods and human factors guidelines for in-vehicle information systems (Green et al., 1991). The objective of this modeling effort was to produce an analytical tool that would serve as an important adjunct to the experiments conducted by UMTRI and would perform one or more of the following functions:

- Provide a basic foundation for understanding driver behavior and performance.
- Design of laboratory simulation and on-road experiments.
- Extrapolation of experimental results to situations not explicitly explored.
- Development and evaluation of guidelines for in-vehicle displays and controls.
- Analytic evaluation of candidate in-vehicle systems.

Data obtained in the laboratory and in on-road studies performed by UMTRI were used to help calibrate and validate the model, and application of the model to the generation of human factors guidelines was demonstrated. (Although models of this type had been previously used successfully in the design of simulation experiments (Levison, 1985), the timing and constraints on the way in which the UMTRI experiments were performed precluded significant application of the model to experiment design.)

Using descriptions of the tasks to be performed and of the driver's information-processing limitations, the model developed in this program predicts a variety of performance measures for typical scenarios. Representative measures include lane deviations, control use times for a variety of in-vehicle systems, and various measures of driver attention such as eye fixations times and scan frequencies, task-to-task transitions, and statistics relating to task interruptions.

The objective of this report is twofold: first, to provide the reader with an understanding of how the model is structured, what kind of input to the model is required to perform model analysis, and what types of output files are produced; and, second, to summarize the results of the model analysis performed in this study. The report is not intended to be a formal users manual.

The material related to model description is organized to accommodate varying levels of interest in the details of theory and implementation. The next section provides an overview intended for readers not familiar with the integrated driver model. Readers satisfied with a top-level understanding of the model may proceed directly to the results presented in chapter 5. Chapters 2 and 3 describe the model structure in more depth; additional details are provided in the appendixes. Chapter 4 describes the information inputs needed to obtain a model solution, as well as the type of data generated by the model. Material on model input appearing in chapters 1 through 3 is summarized for convenient reference, and examples of output data files are presented. Chapter 5 summarizes the model analysis performed in this study, including validation, calibration, and application to guideline development.

OVERVIEW OF THE MODEL

Driving consists of a set of tasks and activities requiring perception, cognition, motor response, planning, and task selection. The latter activity is particularly important as the driver must often choose between attending to the visual scene cues needed for vehicle control and other information sources competing for visual attention (e.g., rear-view mirror, climate control, advanced in-vehicle display, etc.).

These activities are organized in a hierarchical manner. The top-level activity consists of setting overall goals (e.g., drive from point A to point B in the shortest time). A variety of subgoals, or maneuvers, are formulated and satisfied over time in order to achieve these top-level goals. Maneuvers relating to automobile control include the relatively high-level tasks that determine the intended path of the automobile (e.g., pull into traffic, maintain lane position, turn right at the next intersection).

Having defined the maneuvers to be performed, the driver must then elect the lower-level task to be attended to at any given instant. The one task that is always competing for attention is that of vehicle control, which typically consists of maintaining lane position and either speed or headway in such a way that allows the intended maneuver to be carried out safely and efficiently. Other tasks, which may or may not be adjuncts to the vehicle control task, will compete on an intermittent basis at frequencies that vary widely from task to task.

In general, each low-level task includes perceptual (obtain information), cognitive (process and plan), and motor (execute response) components. These processes may be considered to be performed concurrently for some tasks—especially the task of continuous vehicle control. For other tasks, such as reading a message on an advanced in-vehicle monitor and turning it off once read, the perceptual and response activities are separated in time sufficiently to be considered sequential activities.

The model described in this report deals with the lower-level tasks of maintaining lane position while intermittently performing additional in-vehicle monitoring and control tasks. (Speed and headway control were not implemented as part of this effort.) This model, which is called the integrated driver model (IDM), is an integration of two previously existing computerized models referred to below as the procedural model and the driver/vehicle model. The procedural model represents the driver of the vehicle in terms of perceptual, neuromotor, and cognitive responses (Corker, Cramer, and Henry, 1990). Submodels may include visual scanning and detection, auditory perceptual processing, neuromotor reaction time, and choice and decision in the selection of activities. The procedural model deals primarily with in-vehicle auxiliary tasks (i.e., tasks other than continuous vehicle control) and with the task-selection and attention-allocation procedures.

The driver/vehicle model predicts closed-loop continuous control behavior. This model, which is currently used to predict lateral path (steering) control, is based on the optimal control model (OCM) for manually controlled systems (Levison, 1989). The structure and predictive value of the OCM has been verified via extensive application to laboratory and operational manual control tasks, and the OCM has been applied successfully to the design of manned simulation studies

(Levison, 1985). The driver/vehicle model is currently implemented to simulate a constant-speed steering task.

The resulting integrated model allows one to predict continuous steering performance as visual attention is intermittently diverted from the roadway to one or more monitoring locations associated with the auxiliary in-vehicle tasks. The model also allows the driver to attend visually to the roadway while simultaneously processing auditory information. Attention-switching and task selection are made on the basis of time-varying priorities that consider, at each decision point, the penalties for tasks not performed. Presentation of auxiliary tasks is controlled, in part, through dependencies on the state of the driving environment as predicted by the model and, in part, through "scripting" (i.e., state-independent time-based occurrence of events defined prior to the model run).

Figure 1 contains a diagram of the integrated driver model showing the principal functional elements of the model and the major communications paths. To make maximum use of previous implementations, the continuous-control driver/vehicle model and the procedural model are implemented as separate processes. Information transmitted from the procedural model to the driver model consists largely of attentional requests (e.g., the driver does or does not attend visually to the steering task) and parametric changes (e.g., new system dynamics to reflect changes in road surface conditions, a command to change lanes); information transmitted in the reverse direction relates to vehicle states (including path errors) and the driver's perceptions of these variables.

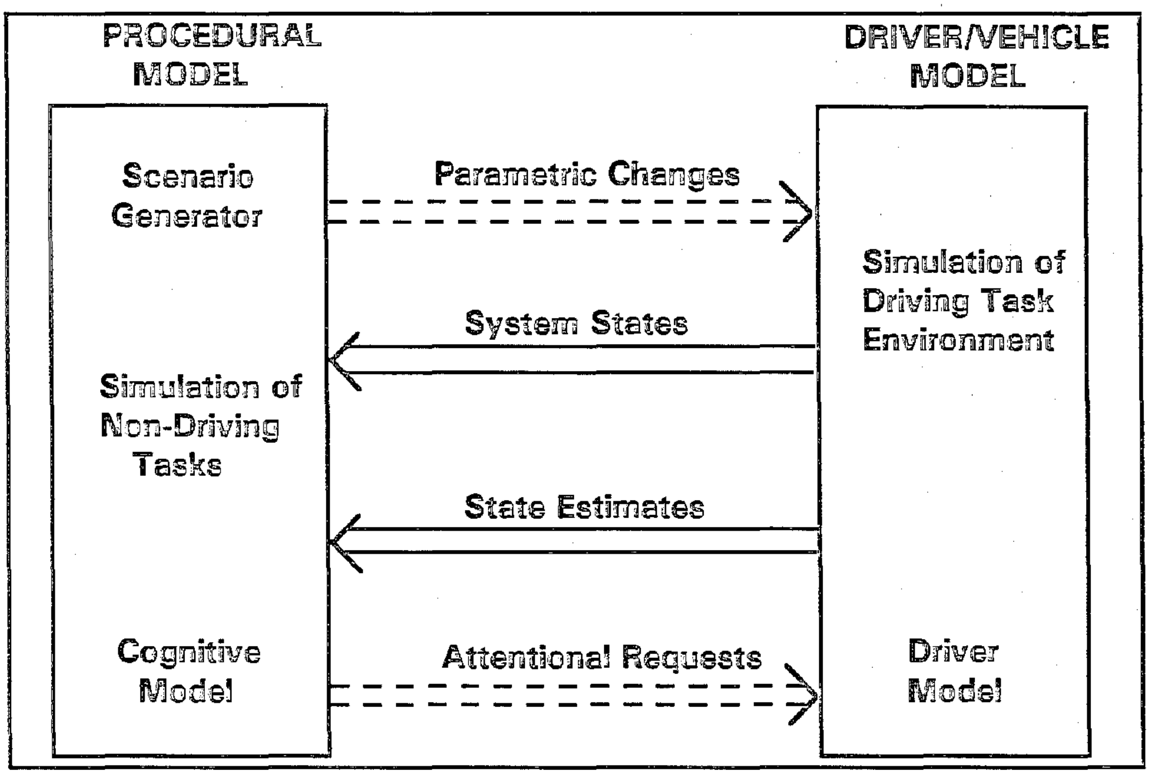

Figure 1. Overview of the integrated driver model.

The Driver/Vehicle Model

The major assumptions underlying the driver/vehicle model are:

- Operator is sufficiently well-trained and motivated to perform in a near-optimal manner subject to system goals and limitations.
- Driver constructs an internalized representation (mental model) of the driving environment in which all dynamic response processes are represented by linear equations of motion.
- Performance objectives can be represented by a quadratic performance index (e.g., minimize a weighted sum of mean-squared lane deviation and mean-squared control activity).
- Driver limitations can be represented as response-bandwidth limitations, time delay, and wide-band "noise" processes to account for information-processing limitations.

In order to obtain a model solution, the user must provide information sufficient to describe the task environment, the performance goals, and the operator's response and information-processing limitations. Because the model is a time-based simulation model, timing parameters must also be specified. The kinds of inputs to be specified for the driver/vehicle model are listed in table 1. These parameters are described in chapters 2 and 6.

Table 1. Inputs to the driver/vehicle model.

Description of Driving Environment
Vehicle response dynamics
Perceptual variables
Command and disturbance inputs
Initial conditions
Driver Characteristics
Mental model of the task environment
Information-processing limitations (S/N)
Perceptual variables (e.g., lane error, heading) used for vehicle control
Time delay
Performance requirements
Motor noise
Simulation Parameters
Simulation update interval
Data recording interval

Predictions of driver/vehicle response are obtained via well-developed mathematical rules for optimal control and estimation (Kleinman, Baron, and Levison, 1970, 1971). Model outputs consist of quantities similar to those measurable in a manned simulation (e.g., time histories for all important system variables), as well as quantities that cannot be directly measured (e.g., the driver's estimate of the value of any system variable), as described in chapter 4 of this report.

The driver's assumed "mental model" of the driving environment is a key feature of the driver/vehicle model. Typically, we assume the driver to be sufficiently well-trained in the specific driving task to allow the mental model to replicate the model of the physical environment. However, we can explore the consequences of the driver's misperception of the external world by making the mental model differ from the world model in terms of parameters values and/or structure.

Further discussion of the operation of the driver/vehicle model component is provided in chapter 2, and implementational details are provided in appendix A.

The Procedural Model

In addition to acting as the supervisory element of the integrated model, the procedural model simulates the in-vehicle auxiliary tasks and performs task selection. We consider first the task selection algorithm and then discuss the overall logic of the procedural model.

Task selection is based on assumptions that are generally consistent with the multiple-resource theories of Wickens and Liu (1988). Specifically, we assume that:

- If two or more tasks require different visual fixation points, only one such task may be performed at any given instant.
- If two or more tasks require auditory input or speech outputs, only one such task may be performed at a given instant.
- If one task requires visual inputs and another requires auditory inputs, they may be performed concurrently (with presumably some performance degradation) if they require different processing codes (i.e., one requires spatial processing, the other verbal processing).
- Task selection is based on the perceived relative importance of competing tasks and is computed by minimizing the expected net penalty of tasks not performed.
- If an auditory and a visual task are performed concurrently, cognitive attention is allocated according to the penalty functions.
- When a task is first attended to, or first re-attended to following attention to another task, attention must remain on this task for some minimum "commit time," after which the driver is free to allocate attention as described above. For a non-interruptible task, the commit time equals the time required to complete the task.

Note that the steering task (which requires attention to the road) is always competing for attention.

The kind of input information to be specified for the procedural model is listed in table 2. The description of auxiliary in-vehicle tasks must be implemented in the computer code, unlike all other procedural and driver model inputs that are specified at run time.

An auxiliary task may consist of one simple activity (e.g., glance at the rear-view mirror) or a sequence of activities, such as the telephone task described later in this paper. An elemental task may require visual attention (eyes), or both visual and manual attention (eyes and hand).

Table 2. Input to the procedural model.

Description of Tasks (hard-coded)
Models of performance versus time
Penalty functions (penalty for not performing task)
Script of Events
Times at which activities are spawned
Simulation Parameters
Simulation update interval

Two categories of parameters need to be specified for each task: parameters relating to performance of the task and parameters that determine the relative importance or urgency of the task. Performance is usually defined by one or more time parameters, which may include (1) times to move eyes and hands, if necessary, in preparation for the task; (2) time to complete the task, and (3) minimum commit time following initial attention to (or resumption of) the task. Some tasks are described by a simple task-completion time. Other tasks are defined by a rate of progress, with the driver allowed to interrupt after some commit time and later continue the task.

For tasks that consist of sequences of activities, sequencing rules need to be implemented, as well as rules for which sequences of tasks must be performed as a unit before the driver is allowed to select another task. (For example, the driver may be assumed to dial the entire area code before deciding whether to look back at the road or to continue dialing.)

Penalty functions for in-vehicle tasks may be specified in terms of a single number, or as a number that (typically) increases with time, up to some limit, until the task is completed. A different kind of penalty function is used for the driving (steering) task; namely, the predicted probability of exceeding a lane boundary within a prediction time that consists of the time required to perform the in-vehicle task segment plus an assumed time to recover control of vehicle path upon re-attending to the road. This computation is based on the driver's current estimate of lane deviation, drift rate, and heading and is similar to the "time-to-line-crossing" metric proposed by Godthelp, Milgram, and Blaauw (1984).

The output file produced by the procedural model includes time histories of the driver's visual fixation point, the position of his free (non-steering) hand, and measures of performance for each in-vehicle task in progress (e.g., number of words read so far from the visual monitor, time elapsed since initiation of the task). As with a manned simulation experiment, post-trial analysis of model outputs can be performed to yield a variety of performance statistics, such as means and standard deviations for all continuous variables relevant to the steering task (including variables internal to the driver), statistics relating to the duration of a given in-vehicle task, and statistics on dwell times and intervals of inattention.

Further discussion of the procedural model component, including descriptions of specific in-vehicle tasks, is provided in chapter 3. Model output is described further in chapter 4.

Simulation Cycle

A simplified top-level diagram of the simulation cycle for the procedural model is given in figure 2. After the model has been initialized, this cycle is executed once per update interval until some stopping criterion has been reached (typically, a stopping time specified at the start of the run). The cycle begins with a check on which new tasks, if any, are to be added to the "active" list (the set of tasks now competing for the driver's attention). New tasks may be spawned according to the time-based script, or because of completion of an antecedent subtask.

If the task currently attended to is locked up, the driver must continue to attend to that task. If the task is not locked up, the task selection algorithm described above is executed to determine the task to be next attended (which may be the same task). Once attention has been determined, all active tasks are, in principle, updated (simulated), not just the task currently attended to. For example, the driving task must be updated each simulation cycle throughout the model run, as the vehicle is always in motion whether or not the driver is attending to roadway cues. Non-attended auxiliary in-vehicle tasks may be updated by having their penalty functions incremented, by having the driver forget information, and so on. Finally, simulation variables needed for post-simulation analysis are recorded in the output file.

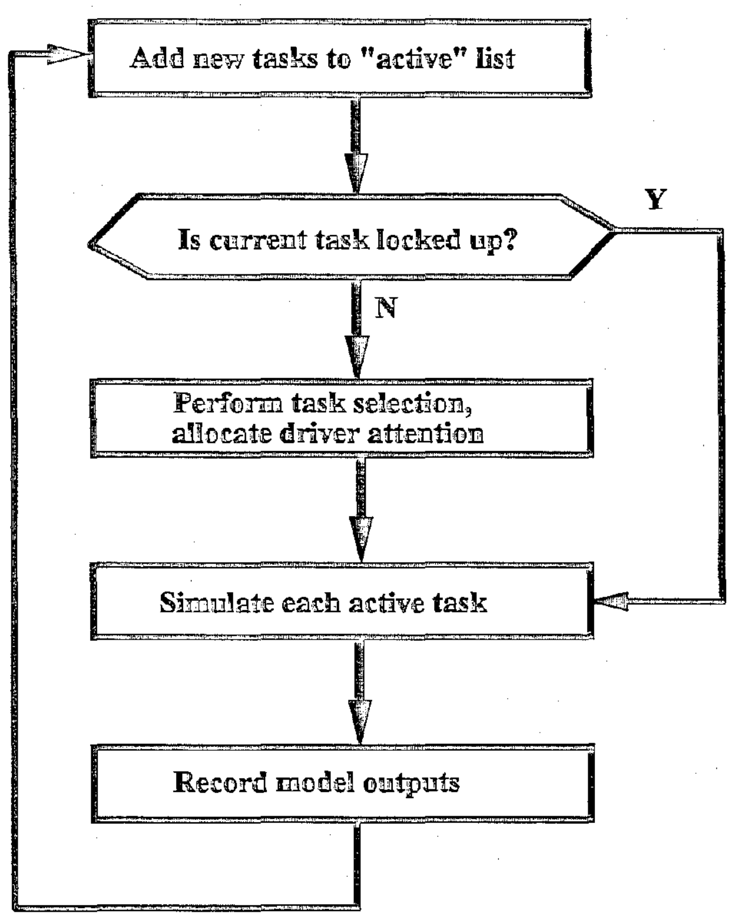

Figure 2. Top-level diagram of the simulation cycle.

2. Theory and Parameterization of the Driver/Vehicle Model

In this chapter we discuss the conceptual driver/vehicle model (i.e., the vehicle control component of the integrated driver model). The discussion focuses on model structure, formulation of the problem, and selection of independent driver-related parameters. The more mathematically intense description of algorithms for obtaining a problem solution are included in the discussion of model implementation in appendix A. The basic assumptions underlying this model have been stated in chapter 7.

We first discuss the three major submodels of the driver/vehicle model: (1) the task environment, (2) the driver's perceptual processes, and (3) the driver's control processes. We then discuss the conceptual integration of these submodels into a unified model.

Submodel of the Task Environment

Both conceptually and implementationally, there are two representations of the task environment to be considered—the representation of the actual or simulated physical environment and the driver's understanding or "internal model" of the task environment. In principle, the physical environment may be implemented to whatever degree of fidelity one wishes, including nonlinear vehicle response components and nonlinear relationships between vehicle states and perceptual cues. (The current model implementation is limited to a linear system representation of the driving task, but this is an implementational convenience rather than a theoretical limitation.) Because of the methods used to compute the driver's estimation and control behavior, however, the internal model must be a linear-system representation of the physical world.

The driver's perceptual and control strategies are based largely on his/her internal model of the task environment; closed-loop driver/vehicle performance is influenced by both the actual and internal models of the task environment. In some cases, these two representations will be identical; in some cases, they will differ. The discussion of the task environment provided in this section pertains to the driver's internal model.

There are two major aspects of the task environment, or problem description, that must be quantified: (1) the "task dynamics," (i.e., the linear equations of motion describing all dynamic relationships among variables not directly related to the driver), and (2) a quadratic performance index or "cost function" to specify the performance goal(s) of the closed-loop system.

Elements included in the description of task dynamics are shown in figure 3. The "plant dynamics" element describes the dynamical response of the real or simulated vehicle to be controlled by the driver. The task environment also includes dynamic response characteristics of external disturbance or command inputs, which are typically modeled as linear processes as discussed below. If the task environment includes panel instruments that have response lags (more appropriate to aerospace operations than to commercial passenger cars), the dynamic equations of motion of such instruments become part of the task description. Also, dynamic response limitations associated directly with the driver's sensory mechanisms (e.g., sensing of whole-body motion via the vestibular apparatus) are included in this element of the task description. (Other aspects of the driver's information processing are included in the submodels associated with driver perception and control.)

One or more external signals acting on the driver/vehicle system must be defined to have a meaningful problem for model analysis; that is, there must be some externally imposed task load. This task load will generally consist of a disturbance force acting to perturb the vehicle from its desired path, or a path command that the driver is to follow. Command inputs may be imposed by the physical environment (e.g., a winding road), or they may be self-imposed by the driver (e.g., a desire to change lanes). Command and disturbance inputs may be applied in combination.

Each stochastic external command or disturbance input (e.g., continuous air turbulence) is modeled as a white noise process that is shaped in frequency by a linear filter and is described by the linear equations of motion describing the filtering process (the "disturbance dynamics" and "command dynamics" blocks of figure 3) and the magnitude of the white noise input.

Continuous inputs that are deterministic in nature (e.g., a specific road curvature profile) are approximated by filtered noise processes that have similar statistical behavior, and these stochastic approximations are included in the driver's internal model. (The model of the physical world, however, will use the deterministic inputs.)

Discrete or transient inputs may be represented (in the driver's model of the task environment) as linear systems with non-zero initial conditions. For some problems, dynamic submodels are defined for each transient input (e.g., a transient gust disturbance profile); in other problems, the transient input may simply be an initial condition on an otherwise existing problem variable (e.g., a lane-change command implemented as an initial error of one lane width imposed on the variable associated with lateral position error).

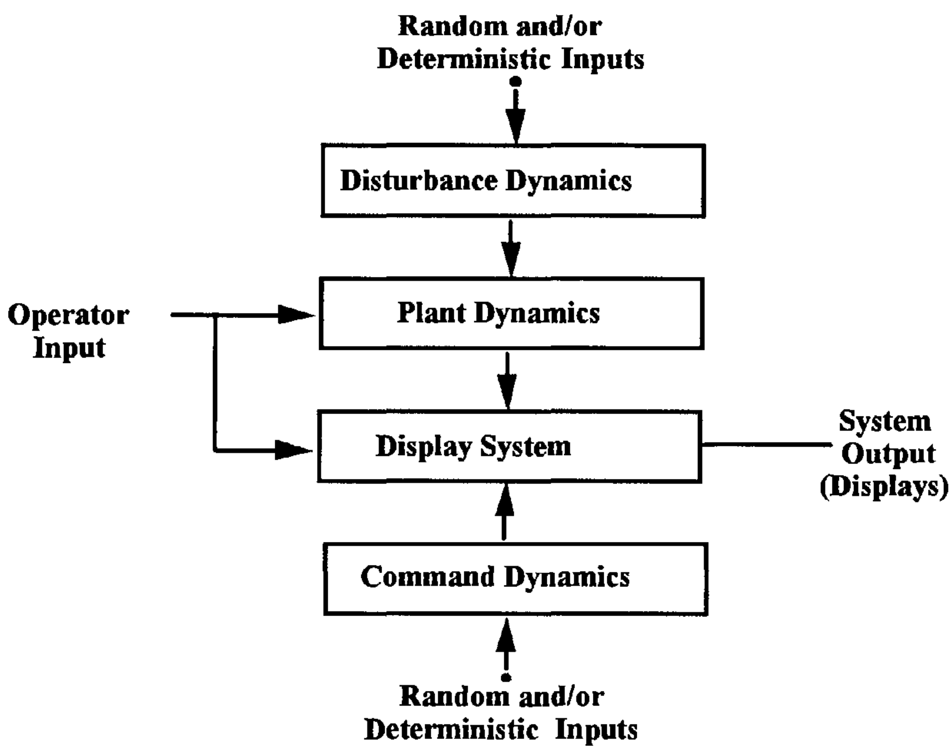

Figure 3. Submodel for the task environment.

As shown in figure 3, the inputs to the task-environment submodel are the external inputs described above, plus the driver's control inputs. The outputs of this submodel are the display variables available to the driver.

Task dynamics (as internalized by the driver) are described by the following linear vector/matrix relationships:

$$\dot{x}(t) = Ax(t) + Bu(t) + Ew(t)$$

$$y(t) = Cx(t) + Du(t) \tag{1}$$

where the vector $\mathbf{x}(t)$ represents the vehicle state, $\mathbf{u}(t)$ is the driver's control input, $\mathbf{w}(t)$ is the white Gaussian noise inputs, $\mathbf{y}(t)$ is the vector of system outputs (displays), and $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, $\mathbf{D}$, and $\mathbf{E}$ are matrices of constant coefficients.

The time-dependent variables included in equation 1 are assumed to be perturbational variables (i.e., variables whose desired values are all zero). If the problem is such that the driver is to maintain one or more variables at a desired state that is not numerically zero (e.g., maintain constant turn rate), the problem is often formulated in such a manner that problem variables are defined in terms of their excursions from the desired state. In this case, a value of zero reflects zero error.

Both conceptually and implementationally, task dynamics are parameterized by specification of the number of state, control, noise, and output variables, and by the numerical entries in the five matrices shown in equation 1. These matrices reflect vehicle, input, and display dynamics integrated into a single vector/matrix equation.

The other aspect of the task environment—task requirements—is expressed by the following scalar quadratic performance index:

$$J = E\left\{\sum_i q_i y_i^2 + \sum_j r_j u_j^2 + \sum_j g_j \dot{u}_j^2\right\} \tag{2}$$

where J is the numerical performance index; E, in this case, is the expectation operator; q_i are the cost weightings associated with display quantities; and r_j and g_j are the weightings associated with the driver's control and rate-of-change of control, respectively. Thus, the driver's assumed performance objective is to minimize a weighted sum of mean-squared display and control quantities, where these quantities represent deviations from desired reference values.

To be consistent with the treatment of task dynamics, one could consider two performance indices —one based on an objective analysis of task requirements and one reflecting the driver's perception of these task requirements. Since only the driver's perception of task requirements influences driver and closed-loop performance, the conceptual and implementational models are concerned with a single performance index, and we classify the performance index as a driver-related independent parameter.

Parameterization of the performance index requires specification of the variables to be included in the computation, plus the associated weighting terms. Application of the optimal control model has led to two philosophies for selecting parameters. In the case of a simple, single-variable, laboratory-type, manual control experiment, the performance index contains two terms—tracking error and control rate. Tracking error is included to reflect the instructions usually given to the subject to minimize mean-squared error and, as we discuss later in this section, control rate is included both to reflect an assumed limitation on the driver's willingness to make rapid wheel movements and to accommodate a mathematical requirement of the problem formulation.

In general, the overall scaling on the performance index is unimportant (it does not influence the optimal driver response behavior), and only the relative weightings are important. Thus, we typically assign a value of unity, and we (iteratively) assign a value to the weighting on control rate to satisfy assumed bandwidth constraints as discussed in the subsequent treatment of the control submodel.

A different philosophy is adopted for more complex tasks. In this case, cost weightings may be derived by first associating a maximum allowable value (or limit) with each variable in the performance index, and then setting the corresponding cost weighting equal to the square of the reciprocal of the limit. Limits are generally determined from considerations of desired performance tolerances, hardware constraints, and driver preferences. This scheme has been used with apparent good results in analytical studies of aircraft operations (Hess and Wheat, 1976; Levison, 1978; Levison and Rickard, 1981).

THE PERCEPTUAL SUBMODEL

Important aspects of the perceptual submodel for the driver are shown in the block diagram of figure 4. The various sources of driver response randomness are lumped into an equivalent "observation noise" process. In effect, we assume that each system output variable utilized by the driver is perturbed by a white Gaussian noise process that is linearly independent of all other such noise processes and of white noises associated with external task demands. Similarly, the pure transport delay associated with human operator response is lumped into an equivalent "time delay" at the perceptual end. All perceptual variables are assumed to be delayed by the same amount.

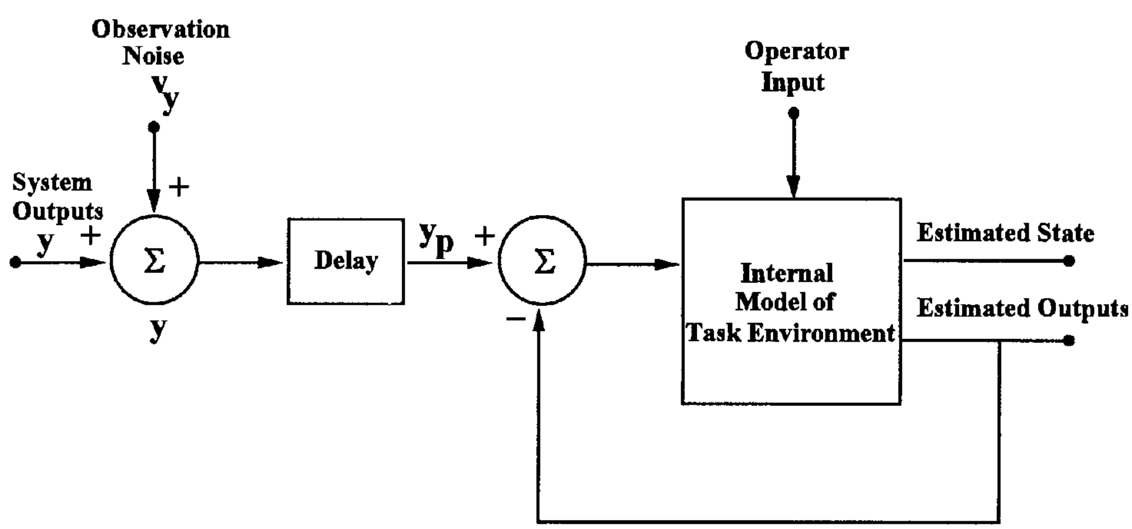

Figure 4. Submodel for the driver's perceptual process.

The association of response randomness and transport delay with the driver's perceptual input is a mathematical assumption rather than an assumption relating to human physiology. From the latter viewpoint, we would assume that these information-processing limitations are distributed throughout the perceptual, information-transformation, and motor response processes. With little loss in generality, however, we gain considerable mathematical convenience by modeling response randomness and transport delay as largely perceptual limitations.

The conceptual model for the driver's perceptual input is expressed as:

$$y_p(t) = y(t-\tau) + v_y(t-\tau) \tag{3}$$

where $y_p(t)$ is the vector of perceived displays, $\mathbf{y}$ $(t-\tau)$ is the display vector delayed by the driver's transport delay τ, and $\mathbf{v_y}$ $(t-\tau)$ is the vector of white observation noise processes. Thus, the driver is assumed to perceive a delayed noisy representation of each display quantity.

Because white observation noise processes are assumed, each such noise process (one per display variable utilized by the driver) is parameterized by a single quantity—the autocovariance (equivalent to π times the power density level). To guide the user (and the computerized models as well) in the selection of appropriate noise covariances, we adopt the following submodel for the observation noise associated with each observational variable.

$$V_y = \pi P(\sigma_y^2 + \sigma_o^2) / f_y \tag{4}$$

where V_y is the noise covariance, π is the scale factor that converts power density level to covariance, P is the observation noise-to-signal ratio, σ_y^2 is the expected squared deviation of the (noise-free) display variable y, σ_0^2 is the variance of a residual noise process associated with perception of the variable y, and f_y is the fraction of attention devoted to y.

A separate relationship of this form is required for each display variable utilized by the driver. In general, the observation noise aspect of the driver model requires specification of an overall noise-to-signal ratio, plus a residual noise variance (or standard deviation) and a fraction of attention for each display variable. (The variable σ_y is computed as part of the problem solution and, therefore, is not a driver-related independent model parameter.)

This submodel is perhaps best appreciated by a review of the way in which it was developed. Early model application revealed that the observation noise covariance tended to scale with the variance of the corresponding display variable for single-axis laboratory tracking tasks utilizing optimal display formats and scalings (Levison, Baron, and Kleinman, 1969). These early results led to the notion of an observation noise/signal ratio, which is consistent with the results of psychophysical experiments that show the human's estimation error variance to scale in rough proportion to the squared magnitude of the physical stimulus.

Subsequent analysis of laboratory tasks requiring concurrent monitoring and control of multiple display variables led to the model for attention in which the observation noise associated with a given perceptual variable scales inversely with the fraction of attention allocated to that variable (Levison, Elkind, and Ward, 1971; Levison, 1979). This representation is consistent with the notion that the operator has a fixed channel capacity that must be shared among the various task-relevant display variables.

Analysis of non-ideal display formats and scalings led to the need to include additional parameters to account for the effects of perceptual resolution limitations (and, in some cases, indifference thresholds representing a minimum acceptable error magnitude). The model implementation used in this study uses the residual-noise variable to reflect limitations of this sort. Additional discussion of the perceptual submodel, including alternative representations for perceptual limitations, are provided by Baron and Levison (1975, 1977).

For simple tracking tasks, the observation noise for <u>all</u> perceptual quantities is specified by a single parameter—the observation noise-to-signal ratio (also known as the observation noise ratio). Typical values for simple tracking tasks are on the order of 0.01 (i.e., - 20 dB). For tasks more representative of actual driving situations, attentional allocations and residual noise will have to be specified as well.

Selection of residual noise levels is highly problem-dependent and often involves a liberal dose of educated guesswork. In principle, each physical display providing task-relevant information to the driver should be analyzed to determine which cues are useful and, for these cues, a quantification of the perceptual limitations. If the display environment is one that has not been modeled previously, the user must attempt to find relevant information from the literature on perception, or conduct a calibration experiment to quantify that aspect of driver performance.

In principle, there are at least two aspects of attention-sharing relevant to the driver/vehicle model: (1) sharing attention between the driving and non-driving tasks, and (2) sharing attention among the perceptual cues used when driving. The latter aspect was bypassed in this study by assuming that the roadway cues were sufficiently integrated so that there were no attention-sharing penalties among these cues (i.e., all f_y were numerically the same). As discussed in appendix A, attention-sharing between driving and non-driving tasks was handled in one of two ways, depending on whether or not the driver was assumed to be able to perform the tasks concurrently.

The remaining driver-related independent model parameter associated with the perceptual submodel is time delay. Values on the order of 0.20 to 0.25 s are typical for driver delay.

In addition to the noise and delay parameters, the perceptual submodel shown in figure 4 embodies the notion of an internal model of the task environment that allows the driver to generate expectations (estimated outputs) of the perceptual inputs. Differences between expected and actual perceptual inputs, along with the driver's perception of his own control input, are used to drive the internal model so as to update the estimated outputs. The internal model is also used to obtain estimates of the (undelayed) system state vector.

THE CONTROL SUBMODEL

The conceptual model of the driver's control activity is diagrammed in figure 5. The block labeled "optimal estimation and prediction," which includes the internal model concept as discussed above, yields estimates of the current vehicle state. These state estimates are then processed by a set of optimal gains to yield a commanded control input. This commanded control is assumed to be perturbed by a motor noise and then filtered by a first-order lag to yield the actual control (driver input) that is applied to the vehicle.

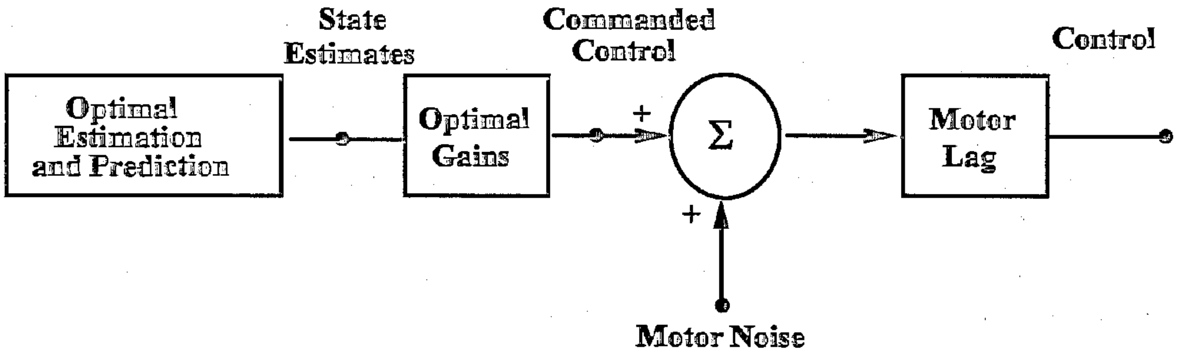

Figure 5. Submodel of the driver's control process.

The motor noise is not intended to reflect a major source of response randomness—that function is fulfilled by the observation noise vector discussed above. Rather, the motor noise is intended largely to reflect limitations on the driver's ability to predict the effects of his control input (e.g., prevent the driver from perfectly predicting the response of the vehicle to his control input).

Conceptually, there are two motor noise parameters for each control input: an actual motor noise and an internal motor noise, each of which is modeled as a white noise process that scales with control-rate variance. The internal motor noise fulfills the function described above. Values of -60 to -50 dB for the internal motor noise/signal ratio have been typically used in previous applications of the optimal control model. As the name implies, this noise influences the driver's response strategy, but it does not represent a physical noise process. (See appendix A for more details on implementation of the internal motor noise.)

The actual motor noise is included for completeness to represent a motor-related noise process that is applied directly to the vehicle. Unless the vehicle control gain is overly sensitive, or the driver has significant neuro-motor deficits, this parameter is usually set to a negligibly small level (e.g., a motor noise/signal ratio of -90 dB).

Early applications of the optimal control model to laboratory tracking tasks suggested that better and more consistent matches to observed manual control response could be obtained by including a performance penalty on the rate of change of control force or displacement, rather than on the magnitude of control force or displacement. A consequence of penalizing control rate is to induce a first-order lag in the driver's control response, the time constant of which is called the "motor time constant." (If the problem contains more than a single control input, this lag is a matrix quantity.) One might also argue from physiological grounds that some form of filtering is required at the motor end to reflect the human operator's inherent bandwidth constraints.

Mathematically, the motor time constant is an aspect of the problem solution and is, therefore, not an independent model input. Conceptually, however, the motor time constant may be treated as an independent-driver parameter for wide-bandwidth laboratory control tasks in which driver response bandwidth—rather than vehicle response bandwidth—is presumed to be the limiting factor. In this case, the penalty associated with control rate is (automatically) iterated until the desired motor time constant is achieved. Motor time constants in the range of 0.09 to 0.12

appear to be typical of wide-bandwidth manual control tasks using force controllers. Larger values are to be expected for manipulators having significant motion.

For relatively low-bandwidth tasks such as vehicle lane regulation, the rate of control is more likely to be limited by physical constraints or by driver preferences, rather than by driver bandwidth. If control activity is limited by physical constraints, the control-rate coefficient is chosen on the basis of a maximum allowable value, as are other cost weightings. The resulting motor time constant then serves as a check that the user has not required an excessive driver response bandwidth. If driver preferences are the determining factor (e.g., the driver wishes to drive in a more relaxed manner), larger values for the motor time constant are selected.

THE DRIVER/VEHICLE MODEL

The conceptual model developed thus far is integrated and summarized in the flow diagram of figure 6. The portion of the driver/vehicle model associated specifically with the driver is enclosed by the dotted lines.

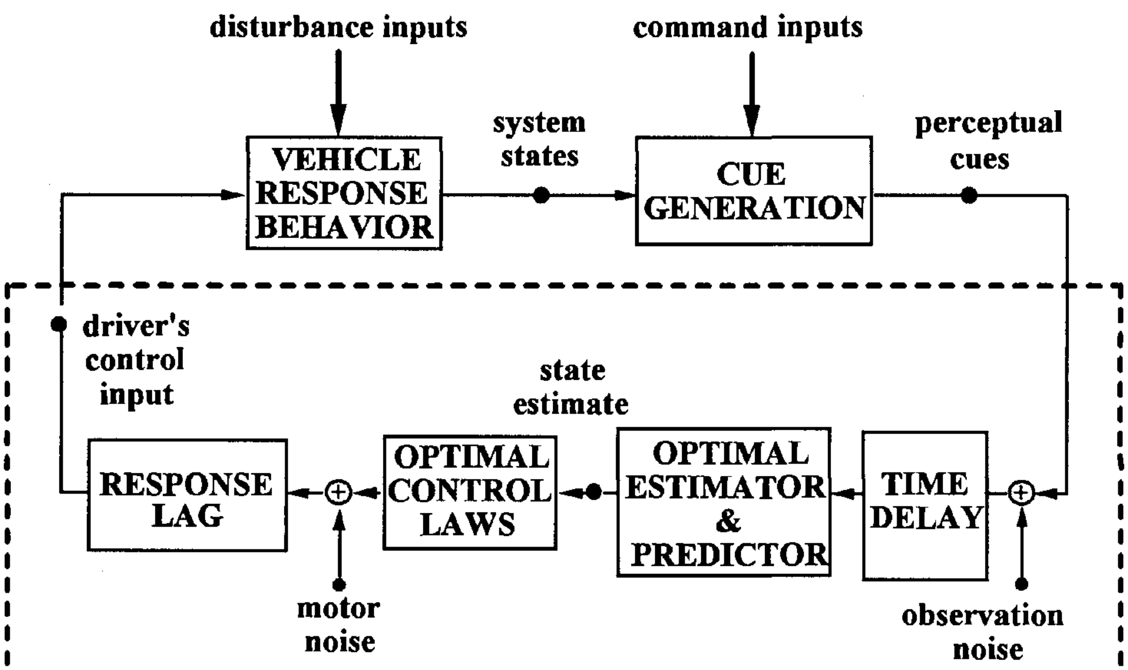

Figure 6. Flow diagram of the driver/vehicle model.

For applications in which the vehicle is maintained at near constant speed and undergoes relatively low lateral accelerations, the model components enclosed in boxes are implemented as linear dynamic processes for which the behavior of the system states is described by a set of linear differential equations. The "Vehicle Response Behavior" element contains a description of the dynamic response of the automobile, the kinematic equations that relate turn rate and speed to lateral displacement, and any dynamic response elements that might be needed to model external disturbances and/or sensor lags.

The "Cue Generation" element accepts the system states and external command inputs to generate the set of perceptual cues assumed to be utilized by the driver. This element contains a linearized approximation that relates the perspective real-world scene cues to system states and command inputs. (For a constant-speed steering task, typical perceptual cues are lane error, drift rate, heading relative to the road, heading rate relative to the road, and road curvature.) These perceptual cues are then corrupted by wide-bandwidth observation noise and delay, where the observation noise reflects both a signal-to-noise type of information-processing limitation as well as perceptual threshold limitations.

The driver's adaptive response behavior is represented by the optimal estimator and predictor, the optimal control laws, and the response lag, with an additional motor noise corrupting the motor response. The "mental model" noted above is a component of the optimal estimator. The estimator and predictor construct a least-squared-error estimate of the current system state, and the (linear) optimal controller generates the optimal control response operating on these state estimates. The motor noise serves to provide some uncertainty concerning the response of the vehicle to the driver's inputs, and the response lag may be thought of as reflecting a penalty for generating a high-bandwidth control response.

The form and quantification of the estimator, predictor, and controller are determined by the specific problem formulation according to well-developed mathematical rules for optimal control and estimation (Kleinman, Baron, and Levison, 1970, 1971). Model outputs consist of quantities similar to those measurable in a manned simulation (e.g., time histories for all important system variables, as well as quantities that cannot be directly measured (e.g., the driver's estimate of the value of any system variable).

When the driver is required to share attention between the vehicle control task and one or more auxiliary tasks (e.g., look at the rear-view mirror, tune the radio), performance of the control task, in general, will degrade. The effects of such interference are accounted for in one of two ways. For intervals in which visual attention is directed away from the roadway cues to some other visual input, the mathematical "driver" receives no perceptual inputs relevant to vehicle control, and the model continues to generate control inputs based on the internal model only. If visual attention is not shortly returned to the roadway cues, the uncertainty associated with the vehicle state grows so large that control effectively ceases.

When the auxiliary task requires listening or speaking, we assume that the driver can attend simultaneously to the vehicle control and auxiliary tasks. In this case, the driver is assumed to continue to fixate on roadway cues, but central-processing resources are now shared between the two tasks. The effects of less than full cognitive attention to the driving task is modeled by degrading the driver's signal-to-noise ratio—in effect, by increasing the observation noise level. (See Levison, 1979, for more details on this aspect of the model.) Either type of attention-sharing tends to decrease the portion of the driver's control response that contributes to effective control (e.g., that decreases steering error) and tends to increase the stochastic component of the driver's control, with the net effect of degrading vehicle control performance.

3. THE PROCEDURAL MODEL

The procedural model component acts as the overall supervisor of the integrated driver model. Specific activities performed by this model component include allocation of attention to the competing tasks, simulation of the auxiliary (non-steering) in-vehicle tasks, and scenario generation (i.e., scheduling of time-dependent events). Task selection and attention-sharing, which was touched upon in chapter 1, is discussed in more detail below. We then describe how the procedural model handles hand and eye movements, which are treated as subsidiary behaviors rather than as explicit in-vehicle activities. Models for simple in-vehicle (non-driving) tasks are described, with descriptions of the more complex tasks explored by UMTRI reserved for the appendix. A discussion of scenario generation rounds out this chapter.

TASK SELECTION AND ATTENTION-SHARING

As noted earlier, task selection is based on assumptions that are generally consistent with multiple-resource theories. Specifically, we assume that only one visual task can be attended to at a given time, only one auditory task can be attended to at a given time, and concurrent visual and auditory tasks may sometimes be attended to simultaneously (with some performance decrement) depending on the cognitive requirements of the tasks. The driver is assumed to share attention in accordance with penalty functions that reflect the consequences of not attending to all the competing tasks. Details of the task selection procedure are provided below.

Assumptions and Constraints

- Except for intervals when the driver is shifting the eye fixation point, total cognitive attention is unity. If the driver is performing a single task, attention to that task is unity. If the driver is performing multiple tasks concurrently, the fractions of cognitive attention devoted to the various tasks sum to unity.
- If two or more tasks require different visual fixation points, only one such task may be performed (attended to) at any given instant.
- If two or more tasks require auditory input or speech output, only one such task may be performed at a given instant.
- If one task requires visual inputs and another requires auditory inputs or speech outputs, they may be performed concurrently (with presumably some performance degradation) if they require different processing codes (i.e., one requires spatial processing, the other verbal processing). If the tasks require the same processing code, then only one may be selected. Thus, the driver is assumed not to be able to read one message and listen to another simultaneously, but he/she can drive and listen (or speak) simultaneously.
- Task selection is based on minimizing the expected net penalty of tasks not performed.
- If an auditory and a visual (typically, driving) task are performed concurrently, cognitive attention is allocated according to the penalty functions.
- If the driver switches tasks, he/she will be committed to ("locked up" by) the new task for some minimum commit time associated with the particular task. If the driver elects to

continue performing the same task, the committed time is some global minimum value that is the same for all tasks. (The simulation update rate of 0.2 s was used in this study.)

- The selection process considers only one step (i.e., the time for which the driver will be committed to his decision); the driver does not plan a series of task-selection maneuvers.

A penalty may be considered to be an index of concern. In the case of a decision to attend to a task other than driving, the driver becomes concerned about the probability of drifting out of the lane, for example. In the case of a monitoring task that is postponed or suspended, the driver becomes concerned about not starting or completing the task. This level of concern may remain constant or may increase with the passage of time. Penalty functions are discussed later in this section.

The task of driving is a special case in the sense that the driver is always driving; that is, we assume that at least one hand remains on the wheel at all times and that some degree of control activity is maintained. The driver does not always obtain visual cues from the roadway scene, however. When attending to the road, the driver is assumed to obtain and use relevant visual cues to update his state estimates. When not attending to the road, no new information is obtained, and the state estimates are updated solely by projecting ahead the current state (see appendix A).

If we wish to model a situation where the driver has but a single explicit task to perform, and performs this task with less than full attention, we must create an artificial task (e.g., sightseeing or daydreaming) that requires the driver to perform the explicit task intermittently or continuously with degraded capability. The rules for sharing attention between the operational and artificial tasks are the same as for sharing attention among operational tasks.

Overview of the Task Selection Procedure

When discussing issues of task selection and attention-sharing, we use the term "driving" to imply attention to the roadway cues, and "auditory" to imply either listening or speaking.

The driving task is always current, i.e., always competing for attention. At any instant, there may be one or more additional visual tasks competing for attention, and one or more auditory tasks competing for attention.

Operation of the task-selection procedure for the most general situation—multiple visual tasks and multiple auditory tasks competing for attention—is diagrammed in figure 7. Net penalties are first computed considering only the competing visual tasks, and the visual task yielding the lowest net penalty is selected for further consideration. Similarly, net penalties are computed considering only the auditory tasks, and an auditory task is selected for further consideration.

Net penalties are again computed, this time comparing the selected visual task with the selected auditory task. If the visual and auditory tasks have the same processing codes (e.g., read text from a monitor, listen to a verbal message), the task associated with the lower penalty is selected for attention. If the tasks have different processing codes (e.g., driving and listening), the driver next performs both tasks, with cognitive attention apportioned according to the numerical values of the penalty functions as described later in this section.

If only visual tasks are current, only the calculation shown in the upper left portion of figure 7 needs to be executed to select the task next attended. If at least one auditory task is current, the competition between visual and auditory tasks must be addressed. (Recall that the visual driving task is always current.)

A special case not shown in figure 7 occurs when a visual verbal task is initially selected, an auditory verbal task is also initially selected, and the auditory task is selected on the basis of subsequent penalty computations. Since we assume that the driver can drive and listen at the same time, we compute penalties as if the driving task has won the visual competition and we apportion cognitive attention accordingly.

Use of Penalty Functions

Penalty functions are computed to determine which task(s) the driver should attend to next whenever two or more tasks are competing for attention and the driver is at a stage where attention can be shifted (i.e., the driver is not committed to continuing the current task). If two tasks are selected for concurrent attention, the penalty functions are also used for apportioning cognitive attention between the two tasks.

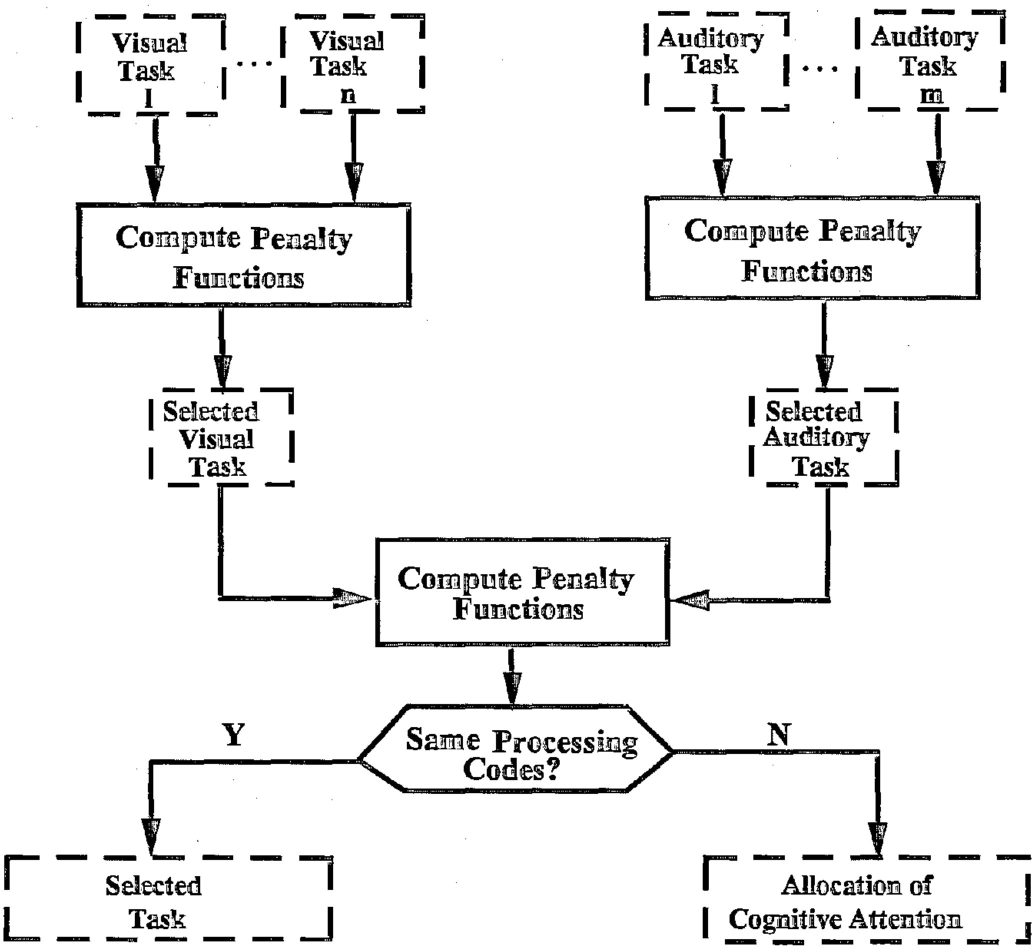

Figure 7. Task selection logic.

The use of penalty functions was selected over the use of expected net gain functions for two reasons. First, for the type of tasks considered in this study, it is generally easier to conceive of a penalty for a task not done than a reward for a task done. (For example, the penalty for not attending to the roadway is an ever increasing probability of colliding with another object or crossing the lane edge; the reward for attending to the roadway is that one minimizes the chance of such a collision.) Second, the penalty function is also more comprehensive than the expected gain function in that it considers interactions among tasks. For example, if more than two tasks compete for attention, the net gain rule would require that we consider each task separately, compute the net gain, and then select the task with the highest gain. On the other hand, if we follow the net penalty rule, the penalty for not performing a particular task, in general, will depend on which alternative is being contemplated.

The operation of the net penalty rule is best explained by example. Assume that the automobile has two in-vehicle auxiliary visual displays: a dedicated navigation display and a display used for providing other information, such as IVSAWS, motorist services, and so forth. Assume that the driver is currently attending to the roadway scene, wants to look at the map on the navigation display, and is alerted to a message on the other monitor. The driver analyzes three options: continue attending to the road, attend to the navigation display, or attend to the other monitor. The net penalty is computed for each option and the option yielding the least net (combined) loss is selected for the next attention interval.

In this example we shall refer to attending to the road as "driving," attending to the navigation display as "navigating," and attending to the other in-vehicle monitor as "monitoring." Penalties are computed as follows:

Option A: continue to drive
Penalty associated with driving =
(penalty for not navigating) + (penalty for not monitoring)

Option B: switch to navigating
Penalty associated with navigating =
(penalty for not driving) + (penalty for not monitoring)

Option C: switch to monitoring
Penalty associated with monitoring =
(penalty for not driving) + (penalty for not navigating)

Thus, the penalty that we associate with a given option is the sum of the penalties for not exercising the other options. In this example, the operator would select the option for which the combined penalty of the other tasks not performed is least. (Later, we discuss a more complicated case in which the driver has the capability to perform two tasks in parallel.)

The penalty computations have two important characteristics. First, they are dynamic variables in that the values computed generally change continuously with time. For example, the penalty for not driving varies with the instantaneous placement of the car within the lane, its drift rate, etc. (Specific penalty functions are discussed below.) The penalty for not monitoring tends to grow with the time elapsed since the monitoring task became current.

The second characteristic is situation dependency; that is, the penalty for not performing a particular task depends on the task currently being performed and the task being contemplated for attention. In this example, there are six possible circumstances in which the penalty for not monitoring might be computed:

Currently driving, considering driving

Currently driving, considering navigating

Currently navigating, considering driving

Currently navigating, considering navigating

Currently monitoring, considering driving

Currently monitoring, considering navigating

In the example discussed so far, all tasks are visual tasks associated with separate look points. Sharing of cognitive attention is not an issue; we assume full attention to whichever task is selected.

The situation becomes more complex if one of the tasks competing for attention is a non-visual auditory (listening or speaking) task, because we assume that in certain circumstances, auditory and visual tasks can be performed in parallel. Let us modify the preceding example by replacing the navigation task by an auditory monitoring task, which we shall refer to as the "listening" task. The three current tasks are characterized as shown in table 3.

Table 3. Characteristics of the three sample tasks.

Task	Modality	Processing Code
Drive	Visual	Spatial
Monitor	Visual	Verbal
Listen	Auditory	Verbal

Task selection and allocation of cognitive attention proceeds as follows. We first consider the competition of visual tasks alone. That is, we compute the penalty associated with driving (i.e., the penalty for not monitoring) and the penalty associated with monitoring (the penalty for not driving). Assume that the driving task is selected on the basis of this analysis. If there were more than one auditory task, a parallel computation would be made, considering only this set of tasks, to select an auditory task for further consideration.

Because the driving and auditory tasks use different processing codes, cognitive attention is shared between the two tasks. An initial allocation of attention is computed as follows:

Attention to driving = (penalty for not driving)/(total penalty)
Attention to listening = (penalty for not listening)/(total penalty)

where the total penalty is the sum of the penalties for not driving and not listening. (Since only two tasks are competing for attention, there is only one computation each for the penalties for not driving and not listening, given the driver's current attentional state.)

Because predicted driving performance is not highly sensitive to relatively modest changes in the allocation of attention between the driving and listening tasks, the allocation is quantized to the nearest 25 percent to provide a smoother-looking time history of attentional allocation. The attention allocated to the driving task is constrained to be at least 25 percent under the assumption that the driver does not completely ignore the driving task in this situation. The competition between the driving and listening tasks thus has four possible outcomes: full, 75-percent, 50-percent, or 25-percent attention to the driving task, with the remainder allocated to the listening task.

Structure of the Penalty Function

In general, a penalty function for a task not performed is conceptually of the form $P=W*f(T,X)$, where P is the penalty, W is a weighting (scale factor), f is some mathematical function, T is a time parameter, and X is a set of relevant world and/or operator states. The time parameter is usually a composite variable that accounts for the time that must elapse until the next opportunity to perform the task. This time will typically include one or more eye movement times plus the time committed to the task contemplated next for attention. If performance of the task requires a hand movement (e.g., the right hand is on the steering wheel and the driver is contemplating a phone-dialing activity), hand-movement time will also be included in the time parameter.

All penalty functions include a scale factor to allow comparison of what might otherwise be incommensurate quantities. It is through the selection of such scale factors that we relate, for example, the undesirability of a 1-percent probability of lane excursion to a second of elapsed time of an uncompleted monitoring task.

The following generic scenarios illustrate how the number of eye movements is determined for a given penalty-function computation. (A similar logic applies to the consideration of hand movements.) Assume that different tasks require different fixation points.

- The driver is currently performing task A and is contemplating next performing task B; we want to compute the penalty for not next performing task A. Two eye movement times must be considered: the time to shift the fixation point from A to B and (presumably the same) time to shift back from B to A.
- The driver is currently performing task A and is contemplating next performing task B; we want to compute the penalty for not next performing task C. Two eye movements are again relevant: the time to shift fixation from A to B and the time to shift from B to C.
- The driver is currently performing task A and is contemplating next performing (continuing on) task A; we want to compute the penalty for not next performing task B. In this case, only one eye movement is considered: the time to shift fixation from A to B.

In the most general formulation of this model for task selection, eye movement times would depend on the location (in terms of visual angle) of the visual fixation points for tasks A, B, and C. To simplify this aspect of the model, however, we assumed in this study that fixation points were either collocated or not collocated. Significant eye movements were assumed not to occur between collocated fixation points (eye movement time = 0), and an average value of 200 ms was assigned to eye movements between fixation points not collocated. No similar assumption was made for hand-movement times, however. Hand movement varied with the distance between manual fixation points as discussed in chapter 3.

Cognitive attention is set to zero when the eyes are in motion. No eye movement time penalty is associated with switching attention to or from a non-visual auditory task.

The penalty function for not driving is defined below. Penalty functions for all other tasks explored in this study, along with values assigned to their independent parameters, are given later in this chapter.

Penalty for Not Driving

The penalty for not driving is defined as the probability, expressed as percent, that the car will exceed the lane boundary T seconds into the future, where T is the prediction interval (defined later in this discussion). A two-step computation is performed. First, the estimated lane position of the car and the uncertainty (standard deviation) about this estimate are computed for T seconds into the future. Under the assumption of Gaussian statistics, these quantities are then used to compute the probability that the automobile will intrude upon either the left or right areas adjacent to the desired travel lane. Unity weighting is assigned to this computation, that is, a unit penalty is associated with a 1-percent chance of exceeding the lane boundary.

Estimated lane deviation is computed as follows:

$$\hat{y}_e\ (T) = \hat{y}_e + T\hat{d}_e + 0.5aT^2\hat{r}_e \tag{5}$$

and the (subjective) standard deviation of the estimation error is:

$$\sigma_{\hat{y}e}\ (T) = [\sigma^2_{\hat{y}e} + T^2\ \sigma^2_{\hat{d}e} + 0.25\ a^2T^4\sigma^2_{\hat{r}e}]^{1/2} \tag{6}$$

where

$\hat{y}\,e(T) =$ estimated path error T seconds into the future

$\hat{y}_e$ = current estimate of path error, feet

$\hat{d}_e$ = current estimate of drift rate, feet/second

$\hat{r}_e$ = current estimate of yaw-rate error, degrees

a = (car velocity in feet/seconds)/57.3, where 57.3 converts radians to degrees

T = prediction interval, seconds

σ = standard deviation of the designated quantity

Note that increasing the prediction interval causes a larger predicted probability of lane exceedance for a given set of path, drift, and yaw-rate errors. Thus, by increasing the prediction interval, we would expect to increase the priority assigned to the driving task in a given situation.

The prediction interval T is computed as

T = eye-movement time + commit time + recover time

The quantity of recover time represents an assumed average time allowed for the driver to arrest a buildup of path error (i.e., nullify the drift rate) that might occur during periods of inattention to the roadway.

The driver model was run to estimate the time to reduce an initial drift rate to 90 percent of its initial value and, on this basis, a recover time of approximately 1 s was derived. This value was treated as a constant for the duration of this study.

As noted above, eye-movement time depends on what task the driver is currently performing and what task the driver is contemplating attending to next. The commit time is similarly situation-dependent. If the driver is contemplating the continuation of the current task, the commit time is simply the minimum for all continued tasks (in this study, the update interval of 200 ms). If the driver is contemplating switching tasks, the commit time used in calculating the penalty for not driving is the commit time associated with the task contemplated for next performance.

As noted below, hand-movement time is not a consideration for the driving task.

CONSIDERATION OF EYE AND HAND MOVEMENTS

Predictions of in-vehicle task performance account for the times taken by any eye or hand movements that are needed to attend to the task. Eye and hand movements are not considered as explicit tasks *per se*. That is, the user need not explicitly include these activities when creating program code for a new in-vehicle task. Rather, the need for eye and hand movements and the times consumed by such actions are determined during the course of the simulation, depending on the current situation.

A display index, along with a physical location associated with this index, are associated with each in-vehicle task and with the driving task. Similarly, a control index and associated physical location are also associated with each task. A display (eye point of regard) location index of 0 is assigned to an in-vehicle activity whenever the eye point of regard is irrelevant (e.g., during a listening task). A control (hand-position) location index of 0 is similarly assigned when hand position is irrelevant. Because only the right hand is free to move, and the left hand is assumed to be always on the steering wheel, the driving task has a control index of 0.

Whenever the integrated driver model simulates a decision to shift attention from one task to the next, the following logic is exercised:

- If the new task has a non-zero display index that is different from the current task's display index, start the timer that accounts for the time consumed by an eye movement.
- If the new task has a non-zero control index that is different from the current task's control index, start the timer that accounts for the time consumed by a hand movement.

At subsequent update intervals, the program checks to see if the times required for both eye and hand movements have elapsed. If so, the in-vehicle activity is initiated. Otherwise, simulation of the activity continues to be deferred.

The times required to make eye and hand movements are determined by whatever models have been implemented by the user to predict movement times. As of the writing of this report, the following models were implemented:

- A fixed eye movement time of 0.2 s.
- The following Fitts' law model for hand-movement time (Card, English, and Burr, 1978):

$$MT = a + b * \log 2(D / W + 0.5) \tag{7}$$

where:

MT = movement time, seconds
a = 0.4 s
b = 0.1 s
W = width of the target (e.g., a pushbutton)
D = distance the hand is required to move

- A set of fixed, empirically based hand-movement times that the user may select at setup time as an option to Fitts' law.

The 0.2-s eye-movement time was selected as a representative value according to data found in the literature (Becker and Fuchs, 1969). Fixed hand-movement times were used for the problems explored in this study. A movement time of 0.2 s was used for the time to move the hand (actually the thumb) from one button to the next for the manually dialed telephone, based on a short experiment performed by the author. Movement times of 0.4 s were used when the hand was moved from one in-vehicle device to the other, based both on data obtained from the literature (Card, Moran, and Newell, 1980) and results obtained from applying the above Fitts' law model to representative situations. (A longer time would be needed to represent movement of the hand from the steering wheel to the in-vehicle device, but since this movement needed to be executed only at the beginning of the model run, this longer movement time was not modeled.)

DESCRIPTIONS OF GENERIC ACTIVITIES

Implementations of relatively simple and generic activities associated with the in-vehicle tasks are described below. Three classes of tasks are described: (1) elemental monitoring and control tasks, (2) reading text from a display, and (3) listening. Complex tasks are, to a large extent, constructed as sequences of these generic activities. The more complex, in-vehicle tasks specifically explored in the UMTRI experiments are described in appendix B.

Elemental Monitoring and Control Tasks

Elemental tasks are those that require a simple operation that is non-interruptible. All such tasks are modeled in terms of time-to-complete. Because these tasks were embedded in more complex, in-vehicle tasks for the situations explored in this study, and because penalty functions were associated with these tasks, penalty functions were not associated specifically with the elemental functions defined here. Completion times were assigned to the elemental functions as shown in table 4.

Table 4. Completion times for elemental tasks with data sources.

Task	Completion Time	Data Source
Observe a simple on-off indicator	0.2 seconds	Teichner and Krebbs (1972)
Read a number	0.15 seconds/digit	Munger, Smith, and Payne (1962)
Read a word for the first time	0.333 seconds	Gibson and Levin (1975)
Re-read a word	0.05 seconds	(author's estimate)
Listen to a word	0.4 seconds/word or as dictated by speech production system	Simpson and Marchionda-Frost (1983) Glanzer, Dorfman, and Kaplan (1981)
Speak a word	0.4 seconds/word or as dictated by the speech recognition system	(same as above)
Push a button	0.2 seconds	(author's estimate)

Reading a Message

The driver's task is to read and understand text presented on a visual display (which may be located either inside or outside the vehicle). This task differs in kind from the elemental monitoring and control tasks described above in that the task may be interruptible, requiring more than a single scan to complete, and a "penalty function" for not attending to this task is defined.

The model of the reading task is based on the following assumptions:

- Depending on the time available for reading, the text may be read in a single scan (glance) or it may require multiple scans, where the intervals between scans are devoted to eye-movements and performance of other tasks (typically driving).
- If multiple scans are required, each scan after the first scan will commence with the driver re-reading, at a relatively high rate, words that have been read on previous scans; the driver will then continue to read new words for the time available at a substantially slower rate.
- The driver is committed to read for some minimum time (typically greater than the base update interval) when a scan is initiated. Subsequent reading during the same scan is interruptible at the basic update interval.

Algorithm for Task Performance

The following variables are defined:

newrate	rate (words/second) for reading new material
nword	number of words in the message
nread	number of words read so far
oldrate	rate for re-reading previously read material
tstart	time this reading task was first initiated
tcommit	minimum commit time when scan is first initiated
tick	simulation update interval
time	current time

All time variables are constrained to be integral multiples of the base tick interval.

When the task first becomes appropriate (e.g., the driver is alerted to the presence of a message), the following initialization is performed:

$$tstart = time$$

$$nread = 0$$

Each time a decision is made to scan (attend to) the display, following attention to the road or some other visual information source, the following calculation is made as to the number of words that will have been read by the time the driver will next have an opportunity to re-allocate attention:

$$nread = nread + (tcommit - (nread/oldrate)) * newrate$$

where *nread/oldrate* is the time required to re-read old material. Thus, the number of words to have been read is the number of words (if any) that have been previously read, plus the number of new words that can be read after the previous words have been re-read. This algorithm assumes that the minimum attention interval *tcommit* is sufficient to re-read previously read words, which was always the case in the situations modeled in this study.

Whenever the decision is made to continue attending to the display, the number of words estimated to have been read by the time of the next decision point (on tick later) is computed as

$$nread = nread + tick * newrate$$

The task is complete and is removed from the list of "appropriate" tasks when *nword* words have been read.

For the situations explored in this study, the variables *newrate* and *oldrate* were assumed to be 3 and 20 words/s, respectively. The minimum commit time was 0.6 s.

Penalty for Not Performing

As discussed earlier, once a task becomes appropriate, a penalty for non-completion is associated with the task and remains in effect until the task is completed. The penalty in effect at any given time, compared to penalties associated with other tasks competing for attention, are used to assign task priorities for the purpose of allocating attention. Penalty functions may assume a number of functional forms, ranging from a simple constant to a value that varies with time in a linear or non-linear manner.

In this study, we assumed that the driver becomes increasingly concerned about not knowing the contents of the message until the level of concern reaches some limiting value. (The limiting value guarantees that there will always be a probability of lane exceedance such that the driving task will take precedence over the reading task.) Specifically, the penalty for not completing a reading task varies linearly with the time elapsed since the task became appropriate, starting with a non-zero initial value and limiting at a maximum value, after which the penalty remains constant until the task has been completed.

The time used in computing the penalty is the time that would elapse between the decision to perform some other task and the next time that the driver would be able to attend to the reading task. This time includes the minimum time consumed in performing the alternative task plus the time, if any, taken for eye movements.

The following variables are used in computing the penalty:

eyemvt	time consumed by eye movements between the monitoring display and some other fixation point
penaltyfactor	rate at which penalty increases with time
maxpenalty	limiting value of the penalty function
minpenalty	initial value of penalty function when task first becomes appropriate
penalty	current value of the penalty function
tmin	minimum commitment to alternative task

Once this task is appropriate, the penalty for not initiating or continuing this task at a given decision point is

$$penalty = minpenalty + penaltyfactor * (time - tstart + tmin(\text{other task}) + eyemvt(\text{other task}))$$

$$\text{if } (penalty > maxpenalty) \quad penalty = maxpenalty$$

The argument (other task) means that the values used for tmin and eyemvt depend on the current eye fixation point and on what other task is being considered for action.

For the situations explored in this study, values of 1.0 unit, 10 units, and 1 unit/s were assigned to *minpenalty, maxpenalty,* and *penaltyfactor*, respectively. As noted above, 0.2 s was assumed to be consumed by each eye movement. An additional weighting term was applied so that 0.2 penalty units would be equivalent to a 1-percent probability of lane exceedance for not attending to the roadway cues. Thus, the relative urgency of the reading task was equivalent to a 0.2-percent probability of lane exceedance when the reading task was first initiated, and it reached a maximum equivalent to a 2-percent chance of lane exceedance.

Unlike most of the parameter selections appearing in table 4, these numbers were not based on published experimental results. Rather, a good deal of "engineering judgment" was used to select values that were expected to produce reasonable predictions of task-sharing behavior. These numbers, as well as the functional forms, are subject to revision if and when experimental data so dictate.

Listening to a Message

The driver's task is to listen to and understand a spoken message provided by an in-vehicle auditory display (message source). The message may be provided by automated equipment (e.g., auditory navigation aid or a vehicle systems warning device), or by people either live or through radio or telephone. Because we do not have a model yet for scanning to and from an auditory display, this model is somewhat simpler than the model described above for the reading task.

The model of this activity is based on the following assumptions:

- The driver must attend to the message continuously from beginning to end to understand it. If the driver "tunes in" late, or interrupts his attention to the message before it is completed, the message is lost and must either then be ignored or replayed if playback capability exists.
- Once initiated, the message is spoken until it is completed.
- Listening is performed concurrently with a visual task, which will generally be that of driving (i.e., looking at the road). Cognitive attention is shared between the listening and visual tasks as described previously in this chapter, except that a minimum cognitive attention of 25 percent is assigned to the driving task when that is the concurrent visual task. No equivalent minimum attention is assigned to the auditory task. Attention to the auditory tasks is considered to be interrupted if the computed fraction of attention to that task falls below 25 percent.

Because we did not locate quantitative models for auditory task performance as a function of attention, auditory tasks were considered primarily as distractions that diverted cognitive attention from the driving task. There was no specific model for how well the auditory task was performed. The only performance-relevant metric was the allocation of attention between roadway cues and the auditory task.

As in the case of the reading task, engineering judgment was applied in the selection of loss function parameter values. The penalty for non-attention to the auditory task was always a constant and was typically set so that the penalty for not attending to the auditory task was equivalent to a projected 1-percent probability of lane exceedance when not attending to the driving task. Thus, at any decision point, if the probability of an anticipated lane exceedance were exactly 1 percent, attention would be split evenly between the driving and auditory tasks. A larger predicted probability of lane exceedance would result in a proportionately greater fraction of attention devoted to driving, whereas a lower probability would allow more attention (up to a maximum of 75 percent) to the auditory task.

SCENARIO GENERATION

The procedural model accommodates a scenario (or script) to specify when time-based independent events are to occur. Typically, the scenario consists of a list of top-level tasks and the times at which each task is to begin competing for attention. Only those tasks that are defined to be spawned at fixed times after the start of the simulation are included in the scenario. Subtasks, or top-level tasks that are contingent upon the completion of another top-level task, are not included the scenario; rather, they are enabled by the precursor tasks as discussed in appendix B.

The scenario does not necessarily specify the times at which the driver begins executing the task —only the earliest times at which execution may begin. The actual start time of a given task depends on the outcome of the task-selection procedure described above.

4. MODEL INPUT AND OUTPUT

The input required for obtaining a model solution and the output produced by the model have been discussed to various levels of detail in the previous chapters. This information is summarized here to provide a convenient reference and a more comprehensive description of model output is provided.

MODEL INPUTS

The inputs consist largely of initialization items that define the problem to be solved and are supplied prior to the model run. In addition, there is input related to the driving scenario, such as road curvature profiles or time-driven discrete events, that is supplied during the course of each run to provide the driver with a non-negligible task workload. This section discusses both types of inputs and provides representative numerical values.

Driver/Vehicle Model

Input for the driver/vehicle model component is categorized as follows: (1) description of the driving environment, (2) driver characteristics, and (3) simulation parameters. Except for command and disturbance inputs, all of the inputs discussed below are of the initialization type and need to be specified only at the beginning of each run, although all but the simulation parameters may be changed one or more times during the course of a model run to reflect changing circumstances.

Description of the Driving Environment

Various inputs are needed to describe the driving environment: (1) system response dynamics, (2) perceptual variables, (3) command and disturbance inputs, (4) performance requirements, and (5) initial conditions.

System Dynamics

Because the integrated driver model (IDM) is implemented to treat only linear system dynamics, there is no need to modify the computer code to handle different vehicles. Rather, the IDM is designed to handle linear systems in a general way, and system dynamics are input by specifying matrices of constant coefficients (i.e., the **A**, **B**, and **E** matrices of equation 1). These matrices must account for all aspects of the driving environment (outside of the driver) described by linear differential equations. System dynamics, at the minimum, will include vehicle dynamics and kinematics. Other linear subsystems, if any, describing the external inputs (e.g., wind, road curvature profile) and hardware or human sensor lags must be incorporated into the overall description of system dynamics.

Disk files containing the constant-coefficient matrices are prepared ahead of time. The appropriate file names are provided at runtime, at which point the dynamics files are loaded into the IDM.

Perceptual Variables

The perceptual variables assumed to be used by the driver must be defined and related to the system state variables through linear transformations. These transformations are embodied in the **C** and **D** matrices of equation 1. (These matrices are included in the same disk files containing the A, B, and E matrices.)

There are two philosophies one might use in defining the perceptual variables: (1) analysis of real-world cues and (2) direct perception of state variables. In the first method, one analyzes the perspective view of the roadway scene, determines what perceptual cues (angles, distances between scene elements, etc.) are likely to be used by the driver, performs trigonometric analysis to relate these assumed cues to vehicle or system states, and then performs a small-signal linearization about the nominal operating point to obtain the desired linear relationships. In the second method, the modeler assumes that the driver directly perceives important state variables. Provided consistent assumptions are made as to how well the driver perceives these cues, the two methods are likely to produce nearly equivalent results. (The quality of the perceived information is discussed in the following pages.)

The second method was followed when performing the model analysis described in the next chapter. In general, the driver was assumed to perceive lateral path error, heading error, and yaw-rate error, where "error" is defined as the difference between the location or orientation of the car and that of the roadway. (All model applications in this study assumed lane-keeping at constant speed.)

Command and Disturbance Inputs

Command and disturbance inputs may be either initialization variables specified at the beginning of the run or continuous time histories supplied throughout the run. To the extent that continuous inputs are modeled as filtered noise (a common way to treat atmospheric disturbances, for example), the filtering portion of the model is embedded in the system dynamics description as discussed above, and a scale factor determining the magnitude of the disturbance is specified at initialization time. In the case of a deterministic input, such as the sinusoidally varying road displacement used in the UMTRI laboratory simulation, samples of the input may be read into the simulation at each update interval. (In the case of the analysis performed in this study, a unit-magnitude waveform was read at each interval and was scaled by a constant scale factor that had been specified at initialization time.)

Initial Conditions

All of the state variables must be initialized at the start of each model run. Typically, state variables represent error quantities as defined above and are initialized to zero, but the user has complete freedom to select initial values.

Driver Characteristics

Driver characteristics are defined in terms of the following parameters: (1) mental model of the task environment, (2) time delay, (3) information-processing limitations, (4) perceptual limitations, (5) motor noise limitations, and (6) subjective performance requirements.

Mental Model

Whereas the system model determines how the vehicle responds to the driver's control input, the driver's response strategy is based on the driver's internal representation of the system dynamics (as well as of other aspects of the driving environment). The internal model is specific in the IDM by defining a set of parameters parallel to those defining the actual driving environment as specified on the previous pages.

A second set of $\mathbb{A}$, $\mathbb{B}$, $\mathbb{C}$, $\mathbb{D}$, and $\mathbb{E}$ matrices are read from the disk file to specify the driver's internal model of system dynamics and perceptual variables. Except as noted below, we generally assume the driver is well-trained in the task under investigation and the internal model is constructed to be identical to the external model. If we wish to explicitly study a situation in which the driver is not fully aware of the current task environment (e.g., effective response dynamics have changed because of a change in road surface conditions), we can do so because some of these matrices will differ between the two models.

The one exception to the general assumption of a veridical internal model concerns the internal representation of a deterministic input (e.g., some specific road curvature profile). If we were to assume a perfect internal model of the deterministic input, the optimal controller would be able to predict perfectly the future course of the external input and would take advantage of this information in generating the driver's response. In theory, a perfectly known input could be followed without the need to look at the road.

In order to represent the human's assumed inability to perfectly predict ahead over a long time interval, the internal model of a deterministic input may consist of a filtered noise process such that the filter bandwidth reflects either the frequency content of the deterministic input and assumed limit on the human's ability to predict ahead, and the amplitude of the noise is chosen so that the magnitude of the filtered noise and the true input are similar. The effect of this approximation is to allow the driver to make short-term predictions of the future course of the deterministic input while still requiring that the driver at least occasionally observe the road.

The mental model of the driving environment also includes the driver's estimate of the initial values for system state variables, as well as a measure of uncertainty about each variable. This uncertainty (i.e., estimation error) is due to an assumed Gaussian random process associated with each estimate that has zero mean and a variance predicted by the optimal estimator. Initial uncertainty is thus specified in terms of the 1-sigma value (i.e., square root of the variance) assumed for the initial estimation error. For tasks in which the input is basically continuous in nature, the effects of the driver's initial conditions will generally "wash out" in 5 to 10 s of simulated time and will have negligible impact on predicted driving performance for the duration of the model run.

Time Delay

The time delay parameter accounts for pure transport delay of information. It is similar to the reaction time measured in psychophysical experiments, but tends to have a larger numerical value. Values typical of closed-loop manual control range from 0.2 to 0.25 s.

Information-Processing Limitations

The driver's overall (cognitive) information-processing capabilities are represented by a parameter called the observation noise ratio—the inverse of a signal-to-noise ratio—that scales the observation (sensor) noise associated with each perceptual variable assumed to be used by the driver. As discussed in chapter 2, the intensive aspect of the driver's attention (i.e., how much the driver is concentrating on the driving task) is reflected through appropriate manipulation of this variable. A typical value of the noise/signal ratio found in laboratory tracking studies is -20 dB relative to the variance of the corresponding perceptual variable.

Perceptual Limitations

Perceptual limitations specific to the ability to utilize (or willingness to control) individual perceptual inputs are reflected by "residual noise" terms associated with each perceptual input. As discussed in chapter 2, this variable may be adjusted to represent either an assumed perceptual resolution limitation or an "indifference threshold" reflecting the driver's indifference to errors below some minimum level of concern.

Motor Noise

Two "motor noise" parameters are defined to account for stochastic aspects of the driver's response associated directly with response execution (as opposed to perceptual or central-processing limitations, which are handled as described above). This parameter is basically a holdover from previous model applications to studies of high-bandwidth control systems having very sensitive responses to control inputs and is not expected to be an important determinant of driving behavior for most model applications. A value of -90 dB (noise variance relative to the squared value of control response) is typically used for the actual motor noise and usually has a negligible effect on problems.

A companion variable is associated with the driver's internal model. Rather than using an internal representation of the motor noise, which was found to induce numerical difficulties, a parameter that we call "relative control uncertainty" is applied to the estimation error associated with the control input. By maintaining a minimum level of control uncertainty, the (mathematical) driver is prevented from making accurate long-term predictions of the effects of his/her control input and, therefore, is forced to pay attention to perceptual inputs.

A value of 0.1 relative uncertainty was adopted for the model analysis performed in this study. That is, an estimation error variance of 0.1 times the squared value of the predicted control input (e.g., wheel deflection) was imposed each time the simulation was updated.

Subjective Performance Requirements

There are two reasons for including performance requirements as an aspect of the driver's characteristics, rather than as part of the description of the driving environment. First, the driver's response strategy is determined by the driver's perception of the performance requirements. Second, performance requirements were not quantitatively specified in the UMTRI experimental program. Rather, performance goals were stated in qualitative terms (e.g., stay in the lane, drive as you normally would, etc.).

As shown in equation 2, the performance index is formulated as the sum of squared values of selected system response and control variables. For the tasks modeled in this study, the performance index was simply a weighted sum of squared path error and squared rate of change of the wheel deflection. The weighting on the lane-error term was adjusted so that unit penalty was associated with a path error that just begins to exceed the lane boundary (a function of lane and car widths). When calibrating the model against experimental data, the weighting on the control term was adjusted to provide a good match to experimental performance measures. When using the model in the predictive mode, the control weighting was either set to a fixed value or adjusted in order to achieve a desired driver response bandwidth.

Simulation Parameters

Two parameters are defined at initialization time to control the simulation: the simulation update interval and the recording interval. A relatively fast update rate is usually desired to minimize the effects of time quantization on the simulation of vehicle dynamics and driver response. To avoid unnecessarily large data files, however, a slower recording rate is selected to be no more than needed to adequately characterize driver and system response behavior. The recording interval is constrained to be an integral multiple of the simulation update rate.

An update interval of 0.05 s was used in this study; recording intervals were either 0.1 or 0.2 s.

Procedural Model

The procedural model differs significantly from the driver/vehicle model in terms of input requirements. Because the driver/vehicle mode is formulated to handle the generic linear control problem, re-coding is not required to handle new driving tasks that fall within this framework, and inputs to this model component are specified at runtime by entering numeric values. To the extent that certain elemental tasks and behaviors are common to a wide variety of tasks, this same approach applies to the procedural model.

The procedural behavior of the driver/operator is captured in an "activity language" (actually, a standardized Lisp form) associated with the procedural model. Underlying this set of activities is the core system responsible for the running of these activities. While this underlying system (or inference engine) remains fixed, the behavior of the operator is independent of the system and this higher level set of activities is input much as one would input data. Extensions to the operator's behavior and/or the specification of new classes of behavior involve the addition of new tasks as described in appendix B.

In this section, we review those inputs that are common to a variety of tasks. The reader is referred to the appendix and to chapter 5 for the parameterization of the more complex tasks explored in the UMTRI experimental program. Inputs are categorized as follows: (1) description of activities, (2) script of events, and (3) simulation parameters.

Description of In-Vehicle Tasks

We review here the independent parameters relating to task performance and penalty computation. The reader is referred to chapter 3 for a detailed discussion of algorithms used to model the in-vehicle tasks.

As discussed in chapter 3 and elsewhere, the allocation of attention among tasks competing for attention is based on the penalty functions computed for the various decision alternatives. The driver allocates attention in a manner that yields the least overall penalty. If visual tasks are competing for attention, one task is given full attention. If a visual and an auditory task are competing for attention and are such that they can be performed concurrently, cognitive attention is apportioned between the two tasks in a manner determined by the penalty computations.

In general, when a task is defined as the concatenation of subtasks, a penalty is associated only with the highest level task. Therefore, no explicit penalty functions are associated with hand and eye movements or with the elemental monitoring and control tasks defined above, as these are generally components of more complex tasks. (These low-level tasks, however, do contribute to the penalties computed for the higher level tasks because the times consumed by these elemental tasks are included in the total times computed for the higher level tasks.)

Hand and Eye Movements

Hand and eye movements are not considered as auxiliary tasks, but rather as subsidiary behaviors potentially required in the performance of certain in-vehicle tasks, depending on the nature of the task and the current eye point of regard and hand position when the driver initiates (or resumes) the task. For the model applications discussed in chapter 5, all eye movement times were considered to take 200 ms, hand (actually thumb) movements required to depress a sequence of buttons for the manually dialed telephone were assumed to take 200 ms each, and other hand movements were assumed to consume 400 ms.

As discussed in chapter 3, the current model implementation also allows the user to specify a Fitts' law model for hand movement as an alternative.

Elemental Monitoring and Control Tasks

Performance on the elemental tasks (monitoring a display for a single-information item, listening, speaking, and pressing a button) was defined in terms of time-to-complete. (Monitoring and control errors were not considered in this study.) Completion times assigned to elemental functions are shown in table 4 and are repeated here as table 5 for convenience.

Table 5. Completion times for elemental tasks.

Task	Completion Time
Observe a simple on-off indicator	0.2 s
Read a number	0.15 s/digit
Read a word for the first time	0.333 s
Re-read a word	0.05 s
Listen to a word	0.4 s/word or as dictated by speech production system
Speak a word	0.4 s/word or as dictated by the speech recognition system
Push a button	0.2 s

Reading Text

Unlike the elemental monitoring tasks, the reading task is interruptible. The driver may read the message in a single glance, or he/she may alternate attention between the in-vehicle display and the roadway cues until the message has been read. Two independent parameters define reading performance: the reading rate for new words and the reading rate for old words. When first attending to a visual monitor, text is read at the rate associated with new words. When re-attending after having interrupted the reading task to attend to the roadway or another in-vehicle display, the driver is assumed to re-read the previously read words at the much greater rate associated with reading old words, then continue reading at the rate defined for new words. For this study, rates of 3 words/second and 20 words/second were associated with new and old words, respectively.

As described in chapter 3, the penalty for not completing the task of reading the text displayed on an in-vehicle display is assumed to increase from some minimum to some maximum level with the time following the instantiation of the reading task. The independent parameters associated with this penalty function and the values used in this study are:

Minimum value	1 units
Maximum value	10 units
Rate of increase of penalty	1 unit/second
Relative weighting	0.2

The "relative weighting" parameter defines the probability of exceeding the lane boundary, expressed in percent, that imposes the same level of concern as one penalty unit computed for the reading task.

Listening or Speaking

There are no independent parameters associated with the performance of auditory tasks of listening or speaking (other than whether or not the task is competing for attention). The penalty function associated with listening and speaking tasks is a constant that is in effect as long as the auditory task is competing for attention. In this study, a value of 1 was typically used, which meant that the penalty for not attending to the auditory task was equivalent to the penalty associated with a 1-percent probability of lane exceedance.

As noted in chapter 3, the parameter values selected for the reading-task and auditory-task loss functions were selected to provide what was felt to be reasonable task-selection behavior.

Driving

Driving performance is computed by the driver/vehicle model component described in chapter 2 and parameterized as described in chapter 4 with regard to performance prediction. The computation of the penalty associated with the decision to not attend to the driving task, however, is performed by the procedural model component.

The reader is referred to chapter 3 for a description of the penalty function associated with the driving task. The penalty computed at any given time is highly state-dependent, and the penalty function consists largely of state variables whose values are predicted by the driver/vehicle model during the course of the simulation This function does, however, contain two independent parameters, which may be modified by the user at runtime: (1) the commit time representing the amount of attention the driver must pay to the roadway cues when the driving task is first attended after an interruption, and (2) a recover time, which is the amount of time we assume that the driver needs to regain effective control of the vehicle after re-attending to the road. Values of 0.4 s and 1.0 s were used for the commit and recover times in this study.

Script of Events

During the initialization phase of each model run, the user has the option to specify one or more in-vehicle tasks that are to be instantiated at pre-determined times during the course of the simulation. Tasks are selected from a menu that contains the names of all tasks that have been coded into the IDM. For each task thus selected, the user specifies the time that the task is to become instantiated. Other parameters needed to describe the task, if any, are also specified (e.g., for a reading task, the number of words to be read). In general, only top-level tasks are specified in this manner; subsidiary tasks performed in the service of the main task will be activated at times determined by the model during the course of the simulation.

Simulation Parameters

This category consists of a single parameter: the update interval for the procedural model component. Because the rate of decision making is assumed to be less than the rate at which control inputs are made, this update interval is generally greater than the simulation update interval specified for the driver/vehicle component. It must, however, be an integral multiple of the latter.

In this study, the update interval for the procedural model was fixed at 0.2 s.

MODEL OUTPUT

Separate, but related, output files are produced by the driver/vehicle and procedural model components. The intent of the primary data file produced by the driver/model component is to facilitate post-trial analysis of time histories of key system variables (e.g., computation of within-trial mean and standard deviation of lane error and wheel deflection, power spectra, etc.). The data file produced by the procedural component is intended more for event-type analysis, such as the computation of scanning statistics.

Sample output files are presented below.

Driver/Vehicle Model

Each model run produces two output files associated with the driver/vehicle component: (1) a log file that documents the inputs to the driver/vehicle model and (2) a data file that contains a frame-by-frame recording of system variables. The log file is a printable (ASCII) file. A binary data file is produced at runtime. This data file may be used directly by programs written to perform post-processing operations on the binary data file, or it may be converted to an ASCII simulation file for visual inspection and for porting to other computers. The log and simulation files are described here.

The Log File

In a typical model run, the log files consist entirely of initialization information defining the model state at the start of the run. Should model parameters be changed during the course of a model run (not including changes in visual attention), these changes and their simulated times of occurrence will be documented in the log file.

A sample log file is shown in figure 8. For purposes of discussion, a "block" of data is defined by text located between blank lines.

The first four blocks of data contain: (1) the name of the log file as stored on the disk, (2) "header" (free-text) information entered by the user to describe the situation being modeled, (3) a record of the simulation update rate and recording intervals, and (4) the name of the pre-recorded file of random numbers used for stochastic inputs, and the first point of that file used in this model run.

The first of the three following blocks gives the name of the files from which the (external) system dynamics were read, followed by descriptive textual information contained in the header of that file. The second block provides the same information for the file containing the system dynamics associated with the driver's internal model. The third block contains the name and header information of a file that provides a textual "wiring diagram" that tells the model how the variables of the external and internal models relate.

The first data block containing numbers in matrix format shows the parameters associated with each of the display (i.e., perceptual) variables defined for the driving task. The first column is the maximum allowable value for each variable (i.e., the value of the excursion from nominal that yields one unit in the performance index). In this example, the performance index consisted of a weighted sum of squared path error and control-rate variation; path error is thus the only display quantity to have a non-zero value. The second column contains the residual noise terms associated with the variables to represent threshold-like perceptual limitations.

Variable names and units are shown for this data block and for others containing variable-specific information.

The final column shows the relative fraction of attention allocated to each display variable. That is, given that some amount of attention is paid to the driving task, these data indicate how that amount of attention is distributed among the variables used by the driver. We assume that the visual scene provides an integrated set of cues that do not compete among themselves for attention. Therefore, variables that are assumed to be used by the driver have a fractional attention of unity; variables not used by the driver have a fractional attention of zero. (The residual noise associated with unattended variables is irrelevant, since such variables are not used by the optimal estimator.)

A similar block of data shows the limit (value for unit cost) associated with each control and control-rate variable. In this example, there is only one control input—the steering wheel deflection. Also shown is the motor time constant. As discussed earlier, the control-rate limit and the motor time constant are closely related; only one of these can be specified as an independent variable. In this example, the zero value indicates to the computer program that the control-rate limit is to be considered the independent variable and the motor time constant is to be computed

```
SAMPLE-1.LOG

Simple lateral dynamics (UMTRI simulator)
Filtered rate control
Simple sinewave input  4.31 ft 0-peak, 3.05 rms (road disp)
Internal model has Butterworth approx to input
Attn to path-err, yrt-err
100 units input = 4.31 ft
ctrl rate limit = 200

Simulation update interval =  0.050 seconds.
Data recording interval =  0.200 seconds.

Randum numbers read from R2000-1.RND
Start at point    1

Linear system dynamics read from file CAR9.LIN
CAR9LIN.WIR
Simple lateral dynamics: lab task
Filtered rate control
Steering gain = 0.048 ft/sec/degree
Input acceleration integrated to yield input rate
Input rate integrated to yield road displacement
ACCINT.DYN
RATEINT.DYN
STEER6.DYN
STRGAIN3.DYN

Operator's internal model read from file CAR9.OPR
CAR9OPR.WIR
Simple lateral dynamics: lab task
Filtered rate control
Butterworth filter approx to sinewave input
Steering gain = 0.048 ft/sec/degree
SINROAD8.DYN
ACCINT.DYN
STEER6.DYN
STRGAIN3.DYN

Input-output linkage read from file CAR9.LNK
Simple lateral dynamics: lab task
Filtered rate control
Butterworth filter approx to sum-of-cosines input
```

Figure 8. Sample log file.

```
-- Display-Related Operator Parameters --
                               Value for   Rms Resid-  Fractional
   Variable     Units          Unit Cost   ual Noise   Attention
1  path-err     feet            4.000E+00   3.000E+00   1.000E+00
2  yrt-err      feet/sec        0.0         3.000E+00   1.000E+00
3  yrt-cmd      feet/sec        0.0         3.000E+00   0.0
4  yaw-rt-err   (dummy)         0.0         0.0         0.0
5  yacc-cmd     feet/sec2       0.0         2.500E-01   0.0

-- Control-Related Operator Parameters --
                              Control      Ctrl-rate   Motor
Variable    Units             "Limit"      "Limit"     Time Const

1  wheel        degrees        0.0          2.000E+02   0.0

- Operator's perception of the external noise SD ---
1  w.n.         w.n.           1.710E-01

Operator delay = 0.375

Motor noise/signal ratio = -90.00
Relative control uncertainty =  0.10
Baseline observation noise/signal ratio = -20.00

 --- Initial values for linear system states ---
Variable    Units         Init Cond
yrt-cmd     feet/sec       1.020E+00
road        feet           0.000E+00
yrt         feet/sec       0.000E+00
path-err    feet           0.000E+00

--- Initial values for operator state variables (internal
model) ---
Variable    Units         Init Cond    Init Ucrty
x1          ***            0.000E+00    1.000E-01
yacc-ipt    feet/sec2      0.000E+00    1.000E-01
yrt-cmd     feet/sec       0.000E+00    1.000E+00
yrt         feet/sec       0.000E+00    1.000E+00
path-err    feet           0.000E+00    1.000E+00
wheeld      degrees        0.000E+00    1.000E-01
wheel       degrees        0.000E+00    1.000E-01
```

Figure 8. Sample log file (Continued).

```
Periodic input read from CURV6.IPT
Period = 26.5 sec
Interval = 0.05
Road: 100 cos(wt+270)
Accel: 5.62 cos(wt+90)

--- Input amplitude  scale factor---
yacc-ipt     feet/sec2     4.310E-02

 --- Initial values for outputs ---
Variable     Units          Init Cond
path-err     feet            0.000E+00
yrt-err      feet/sec       -1.020E+00
yrt-cmd      feet/sec        1.020E+00
yaw-rt-err   (dummy)         0.000E+00
yacc-cmd     feet/sec2       0.000E+00
road         feet            0.000E+00
wheel        degrees         0.000E+00
yacc-ipt     feet/sec2       0.000E+00

Actual control-rate values for unit cost
    2.000E+02

Motor-time-constant matrix
    4.587E-01

--- Simulation terminated at time =  240.000 seconds ---


Data for post-processing saved in file: SAMPLE-1.DAT
```

1 ft = 0.305 m

Figure 8. Sample log file (Continued).

as part of the problem solution. Conversely, a non-zero motor time constant would be interpreted as the independent input and the value for control-rate limit would be iterated upon to yield the desired motor time constant, using the specified value to initialize the iteration.

The following three blocks of data show the driver's internal model of the magnitude of the external command or disturbances input, the assumed time delay, and parameters relating to motor noise and overall information-processing noise/signal ratio. The relatively large time delay shown in this example includes system as well as operator delays.

The next two data blocks show the initial conditions specified for system states, and the operator's internal model of the initial state variables and associated estimation uncertainty (i.e., 1-sigma bound on the estimate).

The following block shows the name of the file containing the road profile along with header information contained in that file; the next block shows the scale factor applied to the road profile.

The block of initial values for outputs is computed by processing the initial state values specified earlier by the display transformation (**C** and **D** matrices). The following two blocks show the control-rate limit and motor time constant actually used in obtaining the model solution. Whichever of these has been interpreted as the independent variable is simply repeated here, and the other shows the value calculated by the model.

The next to last data item shows the time of the next event, which, in this case, was the termination of the model run. Finally, the name of the binary time-history data file is given.

<u>The Simulation File</u>

The initial segment of the simulation file from the same sample model run is shown in figure 9. The first line of the file contains a code word (in this case, "MOD" to signify model output) followed by an integer indicating the number of header lines to follow. (The header will be identical to that appearing in the log file.) The line following the header contains the string "VARIABLES" followed by the number of recorded variables. Each such variable is represented by a line of text specifying its class, name, and units. Use of the class identifier allows multiple use of variable names as demonstrated in the following example.

Five data classes are shown:

1. TIME. This class contains only a single variable—the simulation time (i.e., the number of simulated seconds elapsed since the beginning of the model run).

2. ATTENTION. Two variables are included in this class: visual, the amount of visual attention paid to the driving (steering) task at a given recording time, and cognitive, the fraction of cognitive attention devoted to the driving task. The value for visual is either unity or zero, depending on whether or not the driver was assumed to be attending to the roadway cues. The value for cognitive will be zero if the driver was not attending to the road, unity if the driver was attending only to the driving task, and a value between 0.25 and unity if the driver is sharing attention between the driving task and an auditory task.

3. OUTPUT. This class includes all variables defined as system outputs, which includes perceptual variables utilized by the driver plus other variables of interest (e.g., external inputs and driver control response).

4. OPER-OUT. Shorthand for "operator output," this class includes the driver's current estimates of system outputs. Only those variables serving as potential cues to the driver are included; hence, this class contains fewer variables than the OUTPUT class.

5. OPER-ERR. Shorthand for "operator estimation error," this class contains the estimation errors in terms of 1-sigma bounds on the corresponding variables included in the OPER-OUT class.

Inclusion of the VARIABLES data allows one to identify by name the specific variables to be analyzed by post-processing software.

```
MOD        7
Simple lateral dynamics (UMTRI simulator)
Filtered rate control
Simple sinewave input  4.31 ft 0-peak, 3.05 rms (road disp)
Internal model has Butterworth approx to input
Attn to path-err, yrt-err
100 units input = 4.31 ft
ctrl rate limit = 200
VARIABLES       21
TIME         time        seconds
ATTENTION    visual      fraction
ATTENTION    cognitive   fraction
OUTPUT       path-err    feet
OUTPUT       yrt-err     feet/sec
OUTPUT       yrt-cmd     feet/sec
OUTPUT       yaw-rt-err  (dummy)
OUTPUT       yacc-cmd    feet/sec2
OUTPUT       road        feet
OUTPUT       wheel       degrees
OUTPUT       yacc-ipt    feet/sec2
OPER-OUT     path-err    feet
OPER-OUT     yrt-err     feet/sec
OPER-OUT     yrt-cmd     feet/sec
OPER-OUT     yaw-rt-err  (dummy)
OPER-OUT     yacc-cmd    feet/sec2
OPER-ERR     path-err    feet
OPER-ERR     yrt-err     feet/sec
OPER-ERR     yrt-cmd     feet/sec
OPER-ERR     yaw-rt-err  (dummy)
OPER-ERR     yacc-cmd    feet/sec2
DATA            21    5
     0.00000   1.000E+00   1.000E+00   0.000E+00  -1.020E+00
   1.020E+00   0.000E+00   0.000E+00   0.000E+00   0.000E+00
   0.000E+00   0.000E+00   0.000E+00   0.000E+00   0.000E+00
   0.000E+00   1.000E+00   0.000E+00   1.000E+00   0.000E+00
   1.000E-01
     0.20000   1.000E+00   1.000E+00  -2.040E-01  -1.020E+00
   1.020E+00   0.000E+00  -5.671E-03   2.040E-01   0.000E+00
   8.541E-03  -6.817E-01  -9.995E-01   9.030E-01   0.000E+00
   8.883E-04   7.446E-01   7.918E-01   7.820E-01   0.000E+00
   1.018E-01
     0.40000   1.000E+00   1.000E+00  -4.077E-01  -1.017E+00
   1.017E+00   0.000E+00  -1.714E-02   4.077E-01   5.151E-01
  -2.001E-02  -1.016E+00  -1.875E+00   1.728E+00   0.000E+00
   4.392E-03   6.296E-01   6.759E-01   6.750E-01   0.000E+00
   1.064E-01
   (more data)
                                                    1 ft = 0.305 m
```

Figure 9. Sample simulation file.

Following the information on variables is a line of text containing the string "DATA" followed by integers indicating the number of variables recorded in a data frame and the number of variables recorded per physical record. The remainder of the simulation file consists of frames of data containing numerical values for each of the variables identified above, recorded at the regular recording interval specified during problem initialization.

Procedural Model

The procedural model component produces a file similar in concept to the simulation file produced by the driver/vehicle component in that it records relevant information every time the procedural component is updated. As we shall show, the formats of the two outputs files are quite different. In addition, a time line is displayed on-screen during the running of the model to indicate the instantaneous allocation of attention among tasks.

The On-Screen Display

The attention time line can accommodate up to four tasks. The time line progresses in the manner of a strip chart, with the display moving from left to right, and with time increasing from right to left.

Figure 10 shows a sample time line for the first 9.6 s of a simulation in which a listening task and an in-vehicle monitoring task begin to compete for attention 2 s into the run. (Each vertical division represents 1 s.) This time line displays the following attentional behavior:

- The (simulated) driver devotes full attention to the driving task for the first 2 s.
- Attention is shared between driving (25 percent) and listening (75 percent) for the next 2.2 s (at which point the listening task is completed).
- Attention is time-shared between driving and in-vehicle monitoring for the remainder of the time shown.
- Prior to each re-allocation of attention between the driving and monitoring tasks, the display shows a 0.2-s interval during which attention is not allocated to either of these tasks. This reflects the assumed down time of the perceptual input system during the execution of eye movements.

Data File

Assumptions

In order to understand the structure of the output file, we must first review some of the key assumptions underlying the model as currently implemented. The assumptions listed here reflect assumed limitations on human operator behavior as well as constraints adopted to simplify the modeling task. In this discussion, a "simulation" refers to a single model run.

1. The driving (steering) task is continuous and always competing for visual attention. The driver never completes the driving task before the end of the simulation.

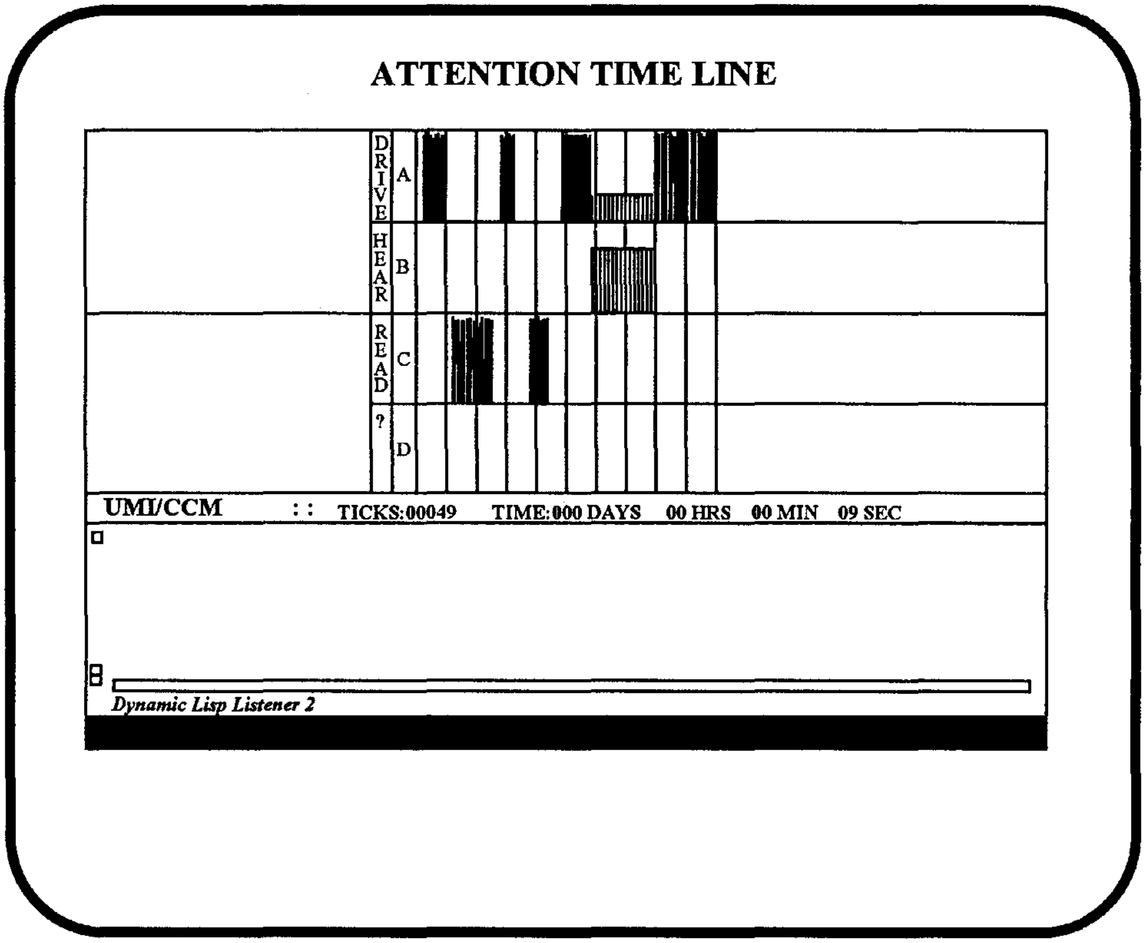

Figure 10. Sample attention time line produced on screen.

2. Any number of visual and auditory tasks may compete for attention. At any given time, the driver can attend to only one visual task and to only one auditory task. Visual and auditory information may be processed concurrently, however, with some performance degradation.

3. Non-driving (secondary) tasks are intermittent in that they have well-defined start and completion times. A given task may occur more than once during a simulation.

4. A non-driving task is typically a sequence of subtasks that collectively define a "top-level task." For example, the top-level task "use telephone" may comprise subtasks that require the driver to: (a) monitor a visual display for instruction, (b) place the call as instructed, (c) verify that the call was placed properly, (d) carry on a conversation, etc.

5. Each subtask requiring visual attention may be considered to have a single fixation point, even though scanning may occur within the display. Thus, when obtaining visual cues relevant to the steering task, the driver is looking at the road, which is typically identified as eye position #1. The fixation point associated with the in-vehicle display providing IVSAWS information may be identified as the "monitor" (or monitor #1 if there is more

than one such in-vehicle monitor) and may be associated with eye position #2, and so on. The model allows more than one task to require the same eye fixation point.

6. One hand is assumed to remain on the steering wheel. The driver's free hand (which may be either the right or left hand, depending on the situation) remains on the wheel unless otherwise required to perform a secondary task. As with eye fixation points, a single hand position is associated with a given task requiring manual activity, and multiple tasks may require the same hand position.

7. All eye movements between separated displays consume 0.2 s. (This assumption is strictly for modeling convenience and could be relaxed if necessary to improve the predictive accuracy of the model.)

8. The driver is assumed to apply full cognitive attention to the tasks at hand. If all tasks competing for attention are visual tasks, a cognitive attention of 1.0 will be associated with the task receiving visual attention. If one visual and one auditory task are being performed concurrently, the cognitive attentions associated with the two tasks will sum to 1.0.

Definitions

The following definitions are needed for the subsequent discussion:

- An "activity" is generally equivalent to a top-level task. In the example presented below, the two activities are driving the vehicle and reading an in-vehicle monitor.
- A "current" activity is one that has been initiated and has not yet been completed. Once completed, the activity is not current until it is re-initiated.
- At any instant, an activity is "active" if it is being attended to by the driver; otherwise, it is "interrupted."
- An activity is "new" for the first frame following its initiation; the activity is considered "continuing" thereafter until completed and re-initiated.
- Each time an activity is initiated marks an "occurrence."
- The time elapsed between the initiation of an occurrence and the subsequent completion time is the "duration" associated with that occurrence.

File Structure

Figure 11 shows the initial portion of the data file produced by the procedural model in a simulation of combined driving and in-vehicle monitoring tasks. The simulated monitoring task consisted of a single subtask: read 10 words from a display.

The data file contains a header that consists of four blocks of information:

1. Free-form descriptive information, not used for analysis.
2. Definitions of the eye positions relevant to the simulation.
3. Definitions of relevant hand positions.

4. Definitions of relevant top-level tasks and subtasks. In this example, each of the two top-level tasks (drive vehicle, read monitor) has but a single subtask (drive vehicle-internal, read monitor-internal).

```
UMI          3
Sample Data File
Four tasks:  UMI:DRIVE-VEHICLE, UMI:DRIVE-VEHICLE-INTERNAL,
   UMI:READ-MONITOR, and UMI:READ-MONITOR-INTERNAL.
EYEPOSITIONS 2
1  ROAD
2  MONITOR
HANDPOSITIONS     1
1  WHEEL
TASKINFO     4
1  DRIVE-VEHICLE
2  DRIVE-VEHICLE-INTERNAL
3  READ-MONITOR
4  READ-MONITOR-INTERNAL
NEWFRAME     1
0.00         1    1
1  2    0    2          1.00 1     1     0
NEWFRAME     1
0.20         1    1
1  2    0    3          1.00 1     1     0
3  4    0    0          0.00 2     1     0
NEWFRAME     2
0.60         0    1
1  2    0    1          0.00 1     1     0
3  4    0    3          1.00 2     1     0
NEWFRAME     2
0.80         2    1
1  2    0    1          0.00 1     1     0
3  4    0    3          1.00 2     1     0
NEWFRAME     2
1.00         2    1
1  2    0         1     0.00 1     1     0
3  4    0    3          1.00 2     1     0
NEWFRAME     2
1.20         2
1  2    0    1          0.00 1     1     0
3  4    0    3          1.00 2     1     0
```

Figure 11. Initial portion of a data file generated by the procedural model.

The remainder of the data file consists of blocks of data describing task status, operator attention, and (optionally) task performance for each frame for which data are recorded. In this example, data have been recorded every 0.2 s—the update rate of the procedural model.

The first record of each data frame contains the keyword NEWFRAME, plus an integer indicating the number of current activities. The second line of information contains the simulation time, an index showing the predicted driver eye position, and an index showing predicted hand position. A value of 0 indicates that the eye or hand is in motion for the subsequent simulation interval. (For example, the zero shown for eye position at simulation time = 0.6 implies that the eyes are in motion between times 0.6 and 0.8.)

The remainder of data for each frame contains a record (typically, one physical line) for each current activity. The records shown in this example contain:

1. Top-level task index.
2. Subtask index.
3. A zero to indicate the end of the task/subtask chain.
4. A status flag (0 = new, interrupted; 1 = continuing, interrupted; 2 = new, active; 3 = continuing, active).
5. The cognitive attention devoted to the activity (0.0 to 1.0).
6. The eye fixation index for the display associated with the subtask.
7. The hand fixation point associated with the subtask.
8. The number of task-specific performance measures following.

A number of task-specific performance measures are potentially available for recording. For example, one could include vehicle path error for the driving task, number of words read so far for a simple monitoring task, duration so far for a telephone conversation task, etc. In the example of figure 11, no performance measures (other than attentional behavior) are shown for any of the activities, and item "8" is 0 for all records.

5. Model Analysis

The first two sections of this chapter are devoted to validation and calibration of the integrated driver model (IDM), where model predictions are tested against data obtained in this and previous experimental programs. The third section demonstrates use of the model to extrapolate the experimental results obtained in this study to situations not specifically explored, but of direct concern to the potential use of in-vehicle automotive information systems. The final section summarizes the foregoing analysis and offers guidance in the use of the model for generating performance predictions.

"Validation" implies demonstration of the general correctness of the model structure and of the model's ability to predict performance (or more realistically, performance trends) across task conditions. "Calibration" refers to the selection of driver-related independent model parameters to best reflect the behavior of drivers in a particular type of operational situation. The distinction between the processes of validation and calibration is somewhat blurred, as is the distinction between those operations and model application to situations where there are data to confirm (or not confirm) model predictions. Often, validation and calibration are performed in the same exercise; that is, by providing a good match to experimental data, one both calibrates the model and demonstrates that the model structure is at least sufficiently robust to explain observed behavior. Similarly, application of the model to additional conditions where experimental data exist offers the opportunity for further validation.

Because of the complexity of the model structure and the relative paucity of experimental data, one cannot expect to quantify unambiguously all the details of the model. A larger data base with more extensive analysis would not completely alleviate this problem, because all one can usefully observe is the external behavior of the human. We cannot, for example, measure the internal decision processes, nor can we directly record the driver's estimates of the current values of important system variables. Accordingly, validation is defined here as the ability to replicate observed performance or performance trends with a consistent and definable set of rules for selecting independent model parameters.

Ideally, a fixed set of independent driver-related parameters would provide a good match to the entire available data base. There are a number of reasons why this goal is not realized, particularly when using data obtained in different studies. Not only are conditions generally different across studies, but there is often insufficient information to describe the task situation to the level of detail required by the model; this means that some judgment must be used in specifying the task-related parameters. Driver populations having different levels of skill, training, and motivation may confound comparisons across studies. Finally, even if all these factors are controlled, it is still the case that what we define as driver-related parameters are not totally independent of the task situation. For example, the residual noise parameters associated with the quality of information available from perceptual cues may vary substantially with viewing conditions (e.g., night versus day), but theory and data may not be available to allow these parameters to be quantified without some form of laboratory or on-road calibration experiment.

For these reasons, it is generally necessary to recalibrate the model when applied to a substantially new driving-task environment and to restrict model applications to variations of the baseline situation for which the calibrated parameter values are likely to remain valid.

Validation and calibration of the two major model elements—the driver/vehicle and the procedural models described in chapters 2 and 3, respectively—are discussed in separate sections. As the reader will discover, the validation data vary in the degree to which they are quantitative. In some cases, there will be comparisons of quantitative performance scores generated by experiment and by the model; in other cases, the validation will consist solely of a demonstration of reasonable model behavior.

In the remainder of this document, the word "prediction" is used as a shorthand to signify model-generated data. It applies both to model results obtained in the process of calibrating against experimental data as well as to output obtained when extrapolating to new situations.

Two models for vehicle (automobile) dynamics were used when performing the model analysis reviewed in this chapter. All analysis of UMTRI laboratory tasks employed a highly simplified vehicle model that corresponded to the simulated vehicle dynamics used in the laboratory tasks. These dynamics did not allow for changes in vehicle heading—the driver directly controlled lateral path velocity (Green and Olson, 1989). A pure time delay of 0.1 s was included in the dynamics to reflect effective delays introduced by the steering mechanism. These simulated dynamics are referred to as the path control dynamics.

All other model analysis of laboratory and on-road driving tasks used more realistic dynamics in which the driver controlled the rate-of-change of vehicle heading. A first- order lag having a time constant of 0.15 s was included in this vehicle model to reflect delays and lags arising, in part, from tire dynamics. These simulated dynamics are referred to as the heading control dynamics.

VALIDATION AND CALIBRATION OF THE DRIVER/VEHICLE MODEL

Validation and calibration of the driver/vehicle model in a single-task (driving-only) situation is reviewed in this section. Discussion is organized into three areas: (1) previous validation of the model structure, (2) calibration and validation using the laboratory data obtained in the UMTRI experiments, and (3) calibration and validation using some of the on-road data also obtained in the UMTRI experiments. As discussed below, a thorough validation and calibration of this model component on the basis of the automobile driving data obtained in this program is not possible. We must rely on the results of previous carefully controlled and designed laboratory studies to demonstrate the validity of the model structure and to provide values for some of the driver-related model parameters.

Previous Validation of the Driver/Vehicle Model Structure

In the simplest of manual control situations there are four independent operator-related model parameters to be quantified: (1) time delay, (2) motor time constant, (3) observation noise-to-signal ratio, and (4) motor noise-to-signal ratio. This is true for the ideal situation in which the displays are designed to minimize the effects of threshold-like phenomena, control gains are optimized, vehicle dynamics respond relatively rapidly, subjects are neither very young nor very old, the performance criterion is specified and easily understood, and subjects are trained extensively and with feedback to maximize performance effectiveness. Furthermore, the external forcing function (command or disturbance) is highly specialized and is designed for maximal analytical properties. (See Levison, 1971, for examples of idealized laboratory tracking experiments.)

In a non-ideal situation, such as would apply in most realistic automobile driving tasks, the following conditions are obtained:

- Perceptual thresholds become important, requiring additional driver-related model parameters to be quantified.
- Drivers do not attempt to satisfy a tight performance criterion, which leads to relatively unconstrained behavior that is sometimes difficult to interpret.
- For on-road studies, the external forcing function is generally not measured (or measurable), which further degrades the possibilities for definitive analysis.

We must, therefore, rely on the basic laboratory studies to demonstrate the validity of the model structure and to suggest values for independent model parameters.

The optimal control model, which is largely embedded in the current driver/vehicle model, was found to provide a good match to data obtained from a range of laboratory tracking tasks using a fixed set on independent model parameters. Furthermore, when the independent parameters were allowed to vary to improve the match to individual data sets, the resulting values of these parameters deviated relatively little from their single-best-fit values (Levison, 1983). This powerful result supports the utility of the underlying model structure as a mechanism for explaining observed behavior and provides some justification for extrapolation to other manual control tasks.

This single set of parameter values does not explain all the reported manual control data, however. In particular, the motor time constant parameter has been found to vary with the simulated vehicle dynamics for dynamics that have significant lags (Levison, 1983). For this and other reasons discussed above, one must generally perform a calibration experiment with baseline data obtained from the situation of interest. In the calibration studies described below, we treat the motor time constant (or an equivalent parameter) as the independent variable of the calibration procedure and, where possible, retain the laboratory values for other driver-related parameters.

Calibration Against UMTRI Laboratory Data

Data from four subjects were provided for the baseline driving task performed on the UMTRI driving simulator (Serafin, Wen, Paelke, and Green, 1993). This simulator contained a steering wheel as an input device and a relatively sparse visual scene roughly approximating nighttime viewing conditions. The highly simplified path-control vehicle response dynamics, although not a true representation of automobile steering response behavior, provided a workload representative of a driving task (Green and Olson, 1989). The subject's task was to remain near the middle of the lane while driving at a constant speed on a road having a sinusoidal curvature with a period of about 26.5 s and a zero-peak lateral deviation of about 1.22 m (4 ft). The data base used for model calibration consisted of one 4-min trial from each of four young drivers.

To calibrate the model, we first set the task-related model parameters to reflect the simulator dynamics and roadway curvature. We then selected driver-related model parameters as follows:

- Response delay was set to 0.2 s, based on previous laboratory tracking data.

- The observation noise/signal ratio was set to -20 dB, again based on previous laboratory tracking data.
- Perceptual noise terms were specified to reflect a visual or indifference threshold of 0.305 m (1 ft) and 0.305 m/s (1 ft/s) for path error and path error rate. These numbers were set substantially higher than associated with idealized laboratory tracking displays to reflect the relative difficulty of obtaining precise information from the simulated perspective roadway display.
- The performance index was defined as the weighted sum of mean-squared lane deviation and steering-wheel velocity. The weighting coefficient on lane error was set so that an assumed maximum allowable error of 1.22 m (4 ft) would provide a unit value of cost. The weighting coefficient on wheel velocity served as the adjustable parameter for this analysis.

Within-trial mean and standard deviation (SD) scores for path error and wheel displacement were computed for each of the four experimental trials, and inter-trial means and standard deviations of the SD scores were computed. Model runs were then generated for various values of the bandwidth parameter until an acceptable joint match to the average path error and wheel deflection scores was obtained. Table 6 shows that model predictions matched the experimental SD scores to within two standard deviations (and to within 10 percent).

Figure 12 shows 2-min segments of the wheel-deflection time histories generated by one of the subjects and by the model using the parameters calibrated as described above. Because the human controller is represented as a mathematically linear system plus "noise," and because the road curvature was sinusoidal, the wheel response is expected to consist of a sinusoidal component of the same frequency as the road sinusoid, plus a stochastic disturbance about this sinusoid. Figure 12 shows that this qualitative description appears to fit both the model and the experimental data, which offers additional validating evidence of the model as a predictor of performance trends. The discrepancies between the high-frequency components of the model and experimental time histories reflect the fact that, by definition, the model cannot replicate (other than statistically) the stochastic component of the driver's time history.

Table 6. Comparison of experimental and model SD scores.

Variable	Experiment	Model
Path error (feet)	0.689 (0.072)	0.640
Wheel deflection (degrees)	17.7 (1.15)	19.1

Calibration Against UMTRI On-Road Data

The model was subsequently calibrated against data obtained from on-road studies conducted by UMTRI during the task K phase of the experimental program (Green, Hoekstra, Williams, Wen, and George, 1993). The data used for calibration were obtained from intervals of near-constant-speed travel on interstate highways at times when the subjects were not required to scan the in-

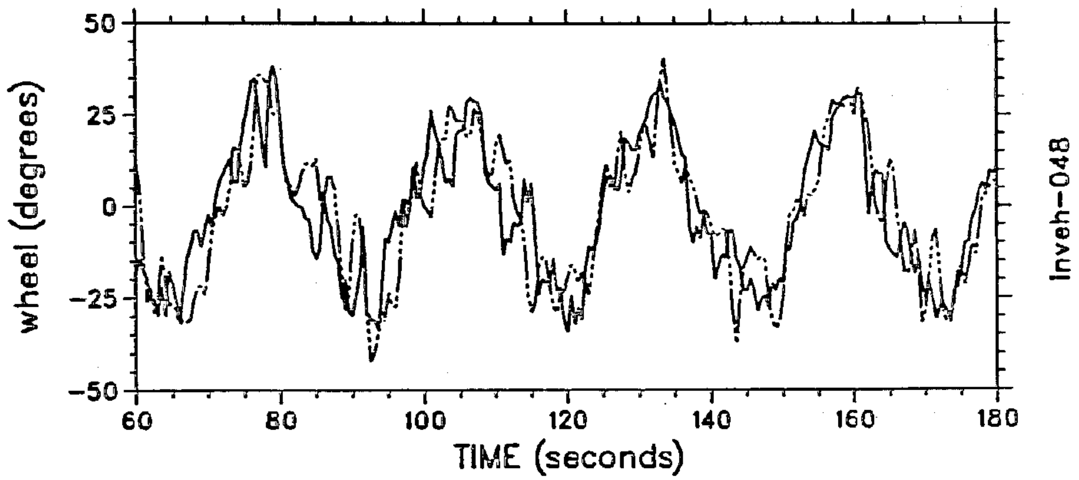

Solid curve: subject. Dot-dash curve: model.

Figure 12. Experimental and predicted wheel time histories.

vehicle displays at frequent intervals (i.e., periods when there was presumed to be little distraction from the driving task).

Measurements of both lane position and steering wheel position obtained in the UMTRI on-road experiments contained a substantial amount of noise (i.e., rapid fluctuations that could not reasonably be attributed to the driver's control activity) that may have been due to errors in data collection hardware and software, including data loss or instability of the lane tracker. Signal conditioning was therefore necessary to obtain reasonable-appearing time history plots and to minimize the effects of this measurement noise on statistics (typically, within-run standard deviations) computed from the time histories.

Procedures for conditioning the signals are described below, followed by the results of the calibration effort.

Signal Conditioning

The need for data smoothing is demonstrated in figure 13, which compares unconditioned with conditioned time histories from one subject driving at around 96.6 km/h (60 mi/h) on a relatively straight section of highway. Lane-position error is shown in figure 13(a) and steering wheel displacement is shown in figure 13(b). Both time histories have been time-shifted to define the start of the analyzed segment (which occurred many minutes into the run) as time zero. Because both recordings contained amplitude offsets, both time histories have been amplitude-shifted so that a value of 0 represents approximately zero lane error or zero wheel deflection.

Two algorithms were applied to condition the data. The first was designed to remove "spikes," where a spike was defined as a data point that deviated by an unusually large amount from the neighboring data points. After some exploration of the algorithm for detecting spikes, a criterion deviation of 1.5 standard deviations was adopted to define a spike.

(a) Path

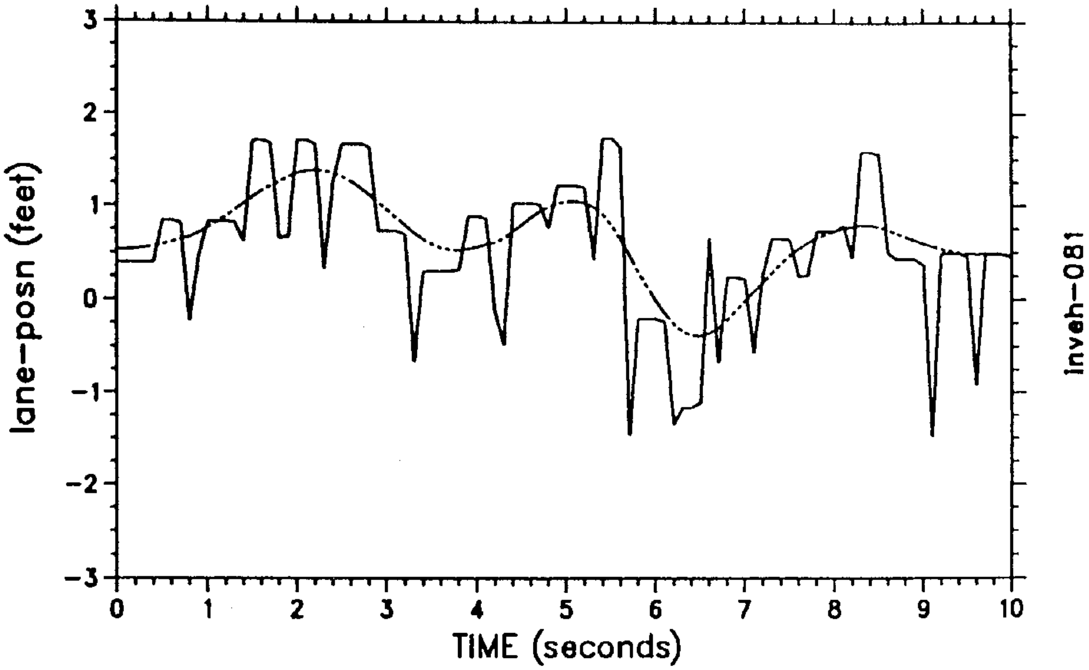

(b) Wheel

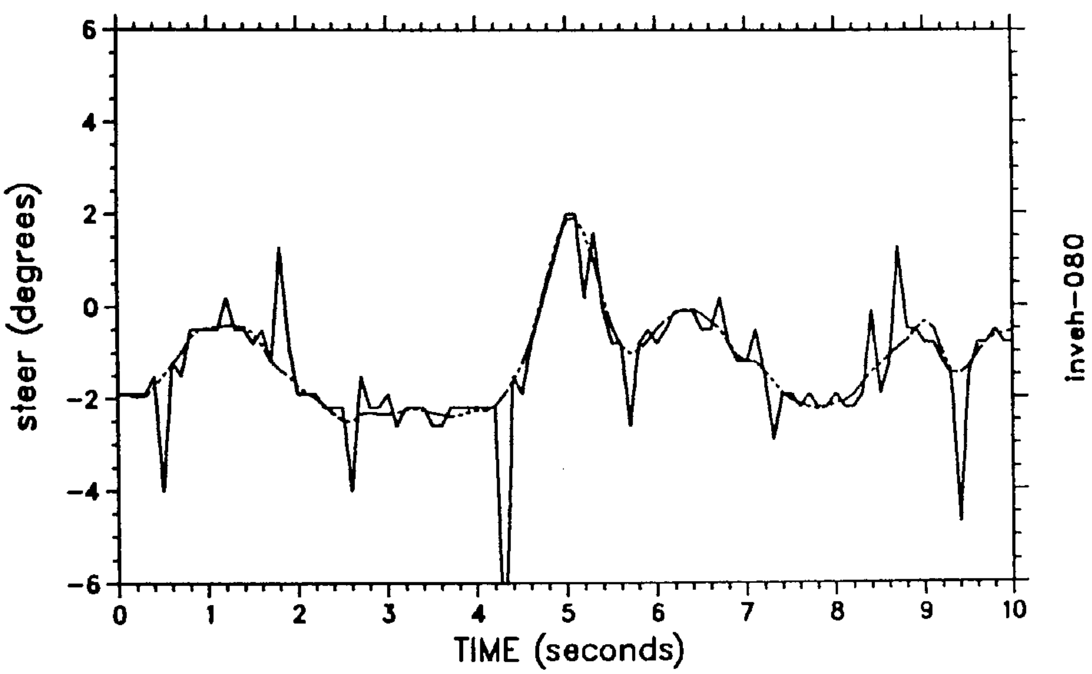

1 ft = 0.305 m

Solid curve = unconditioned data. Dot-dash curve = conditioned data.

Figure 13. Effects of signal conditioning on experimental path and wheel time histories.

The second algorithm consisted of smoothing (filtering) the time history. Because smoothing was performed after all data had been recorded, rather than in real time as the recordings were being obtained, it was possible to adopt a smoothing algorithm that did not introduce a time shift in the smoothed time history. This feature allowed us to use different smoothing time constants for the steering wheel and lane position time histories without influencing the relative timing of the resulting signals.

A "boxcar" smoothing algorithm was applied in which a data point was replaced by the short-term average of the data points contained within a window ranging from some specified time interval preceding the occurrence of the data point to an equal interval following the data point. We designate this interval (one-half the width of the window) as the time constant. After testing of this algorithm against experimental data from a number of subjects, smoothing time constants of 0.1 s and 0.5 s were selected for application to steering wheel and lane position data, respectively.

Smoothing was not performed in place; instead, a second time history was constructed of the smoothed data points. After smoothing was completed, the original time history was replaced by the smoothed time history.

Two passes each of the spike detection and smoothing were applied to each time-history segment analyzed. Spike detection was performed first, then smoothing. The procedure for conditioning the signals was as follows:

1. Compute the mean and standard deviation of the unmodified time history. Detect and remove spikes, replacing each spike with the average of the adjacent data points.
2. Repeat step 1 using the modified time history.
3. Perform the boxcar smoothing on the time history with spikes removed. After smoothing is completed, replace the original time history (minus spikes) by the smoothed time history.
4. Perform a second boxcar smoothing operation on the time history as described in step 3.

Note that performing boxcar smoothing twice is equivalent to performing a "triangular" smoothing in which the weights given to the data points included in the short-term average decrease with the distance of a data point from the center of the averaging window.

The analysis presented so far gives an indication of the noise-reducing power of the signal-conditioning algorithms, but it does not reveal the inevitable distortion imposed upon the noise-free portion of the signal. Because noise-free data were not available, model analysis was performed to estimate the degree of signal distortion. Figure 14 provides a visual indication of the waveform distortion expected from the smoothing process in the case where the driver causes a rapid shift in lane position. Shown in this figure are the unconditioned and conditioned time histories representing vehicle lateral path response (figure 14(a)) and steering input (figure 14(b)). The signal conditioning algorithms described above were applied. Because these time histories were entirely analytic, there were no spikes, and only the smoothing algorithms affected the waveforms.

A single cycle of a sine wave was used to represent the steering input. A period of 1 s was adopted to reflect a maximally rapid steering input anticipated from a human driver, based on inspection of UMTRI driver data. This input was applied starting 1.5 s into the model run. The

(a) Path

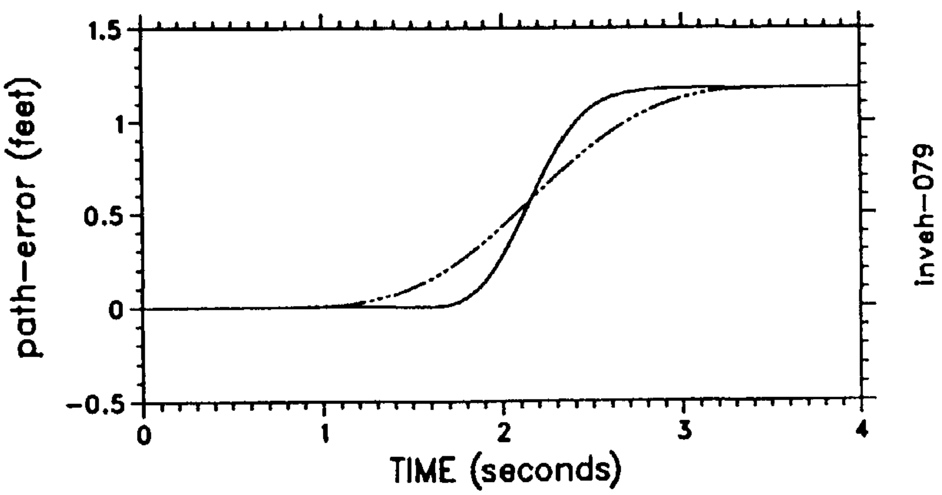

(b) Wheel

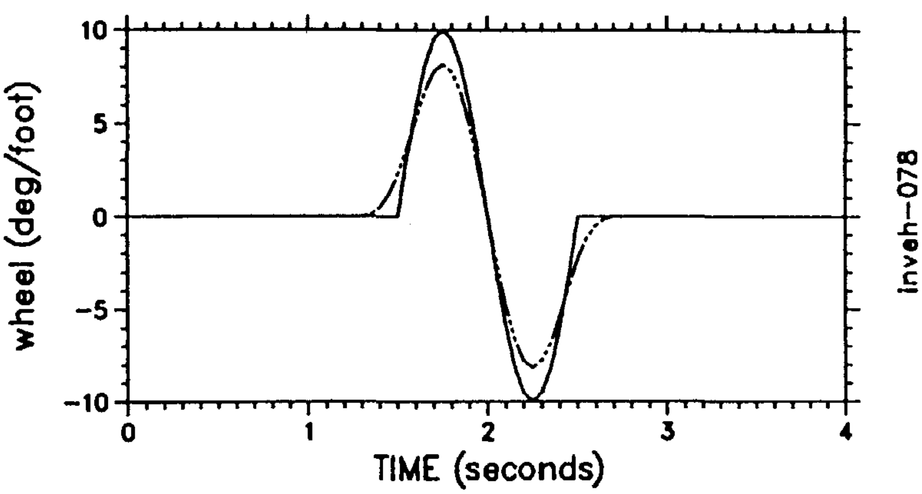

1 ft = 0.305 m

Solid curve = unconditioned data. Dot-dash = conditioned data.

Figure 14. Effects of signal conditioning on predicted lane-position maneuver.

vehicle response was obtained by applying this sinusoidal pulse as the input to the linear dynamic model formulated to approximate the response of the vehicle used in the UMTRI on-road experiments.

Figure 14(a) shows that the steering pulse was spread out in time by a total of about 0.4 s by the smoothing process, and the peak amplitude was decreased by about 20 percent. The amplitude of the resulting lane deviation was unaffected by smoothing, but the rise time—defined here as the time interval between 10-percent and 90-percent response—increased from about 0.6 s to 1.3 s.

While this analysis gives an indication as to how the response to a rapid input might be influenced by the signal conditioning described above, it tends to exaggerate the effects of filtering with respect to data generated by the UMTRI experiments, where sudden changes in lane position were not generally anticipated. To provide a more relevant estimate of how analysis of the UMTRI driving data were likely to be affected, the integrated driver model was used to generate continuous steering and lane-error time histories under conditions representative of the baseline (straight highway) driving situation. Independent driver-related model parameters were selected to match wheel and lane-error statistics as described in the next few pages. A simulated 4-min driving task was analyzed in which residual disturbances such as road-surface irregularities, wheel play, etc., were assumed to provide a high-bandwidth disturbance to vehicle yaw rate. The 4-min data segment was preceded by a 10-s warm-up interval to allow initial transients to wash out.

Figure 15 compares conditioned with unconditioned path-error and wheel time histories for a 10-s segment. As expected, signal conditioning somewhat reduces the peaks of the short-term fluctuations.

Within-trial standard deviation (SD) scores for path-error and wheel time histories were computed for the entire 4-min data period for both unconditioned and conditioned data. Table 7 shows that signal conditioning reduced the predicted path-error SD score by about 4 percent and the predicted wheel SD score by about 14 percent.

Table 7. Effects of signal conditioning on predicted SD scores.

	Unconditioned	Conditioned
Path error	0.705	0.674
Wheel	1.054	0.904

The signal conditioning algorithms had a proportionately greater influence on the predicted wheel score than on the path-error score, even though more smoothing (i.e., larger time constant) was applied to the path data. This seemingly contradictory trend is explained by the frequency composition of the path and wheel signals. Lane position tends to exhibit a significant amount of low-frequency drift (see figure 15(a)), whereas wheel deflection cannot contain low-frequency drift if the car is to remain on the highway. As a result, the path time history contains a substantially larger fraction of its signal power at relatively low frequencies, where the effects of smoothing are substantially less.

(a) Path

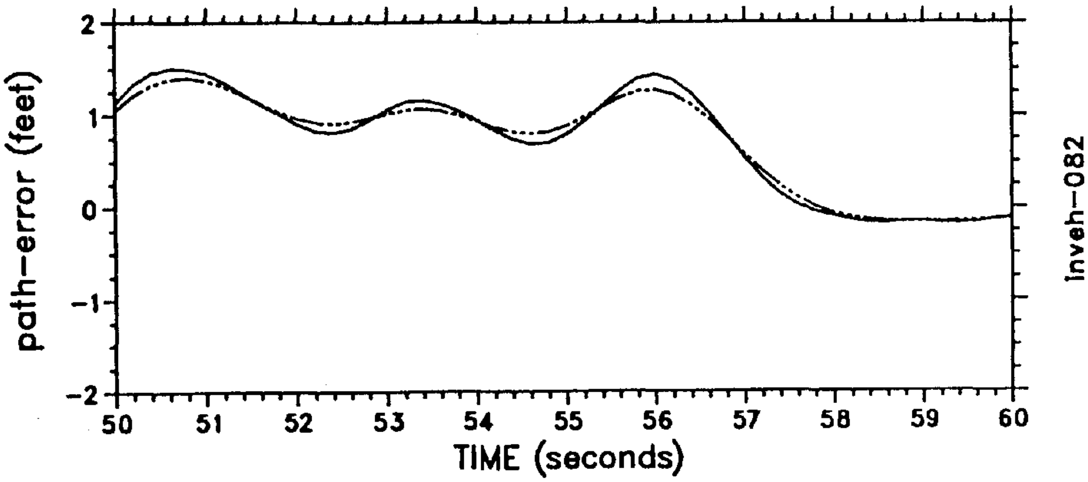

(b) Wheel

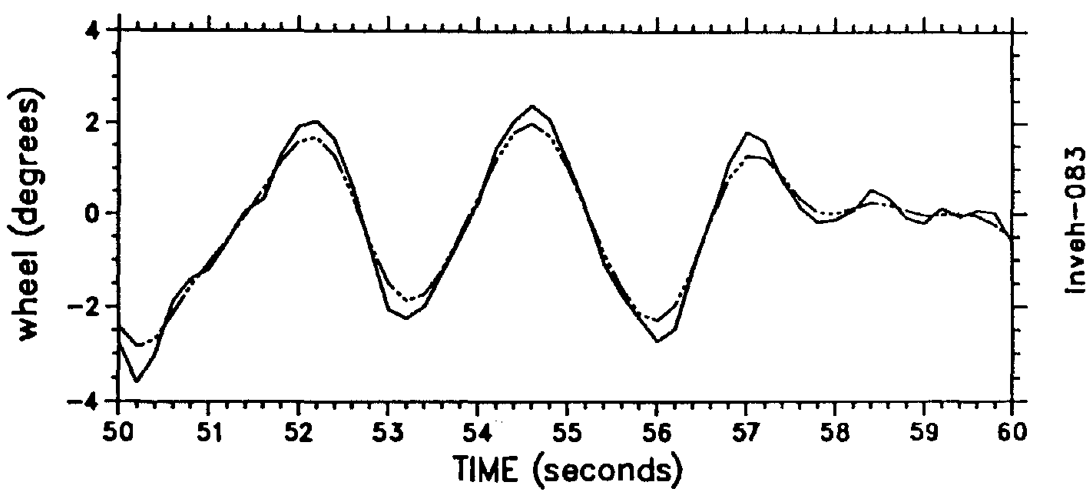

1 ft = 0.305 m

Solid curve = unconditioned data. Dot-dash curve = conditioned data.

Figure 15. Effects of signal conditioning on predicted path and wheel time histories.

Calibration Against Baseline Data

The driver/vehicle model component was calibrated for application to the on-road experiments with the aid of baseline data obtained in task K of the UMTRI experimental program (Green, Hoekstra, Williams, Wen, and George, 1993). As mentioned earlier, the model was tested against path-error and wheel data obtained during intervals of near-constant-speed travel on interstate highways at times when there was presumed to be little distraction from the driving task. Average within-trial SD scores were computed from data provided by 25 subjects. This pool included old and young, and male and female drivers. The heading-control vehicle model was used for this analysis.

As noted above, the simulated 4-min driving task was analyzed in which residual disturbances such as road-surface irregularities, wheel play, etc., were assumed to provide a high-bandwidth disturbance to vehicle yaw rate. The driver's task was assumed to be that of keeping the car near the center of the lane. The 4-min data segment was preceded by a 10-s warm-up interval to allow initial transients to wash out.

This calibration exercise differed from the one based on the laboratory data in that the characteristics of the (assumed) external disturbance source(s) were unknown. It was therefore necessary to calibrate the assumed external disturbance as well as the driver model. The external disturbance process was modeled as a stochastic noise process acting on the orientation of the front wheels. (See Zwahlen and Balasubramanian, 1974, for analysis and application of a similar model.) This noise process was assumed to have a bandwidth substantially greater than the closed-loop bandwidth of the driver/vehicle system. The magnitude of this noise process was determined through the calibration procedure.

The model was calibrated as follows:

- All driver-related independent model parameters, other than the coefficients associated with the performance index, were held constant at the same values used to calibrate against the laboratory data.
- In addition to lane position and heading, the drivers were assumed to use vehicle yaw (turn) rate as a perceptual cue, and residual noise was adjusted to reflect an effective perceptual threshold of 1 deg/s.
- The cost-weighing coefficient on lane error was selected to correspond to a maximum allowable error of 0.914 m (3 ft). This represented the maximum lane deviation that a 1.83-m (6-ft) wide car could undergo within a 3.66-m (12-ft) wide lane without crossing the lane boundary.
- The motor time constant parameter was treated as a dependent variable of the calibration procedure. Because of the coupling between this parameter and the cost coefficient associated with steering wheel rate of movement (i.e., steering wheel velocity), varying this parameter was functionally equivalent to varying the cost coefficient.
- The magnitude of the external disturbance process was also treated as a dependent variable of the calibration procedure.

The model was first calibrated against the data corresponding to a nominal 96-km/h (60- mi/h) forward speed. Round-number values of motor time constant and of the external noise process were explored to yield the closest joint match to experiment lane and wheel SD scores. A subjectively good match to the scores was obtained with a motor time constant of 0.25 s (corresponding to a wheel velocity maximum allowable value of 32.7 deg/s) and a noise covariance of 0.002. The consequences of this particular noise value in terms of its effect on the vehicle's trajectory will be demonstrated shortly.

To test the predictive usefulness of this calibration, the model was then used to predict the scores corresponding to a forward speed of 64 km/h (40 mi/h). Because the relationship between motor time constant and cost coefficient is dependent on vehicle dynamics, two predictions were made: one keeping the motor time constant the same as in the previous calibration run and one keeping the wheel velocity cost coefficient the same. The disturbance noise covariance was not changed. The better match to the 64-km/h data was obtained with a fixed wheel-velocity cost coefficient, and this cost coefficient was used in subsequent model analysis.

Measured and predicted SD scores are compared in table 8. Predicted scores were computed from model-generated time histories that were filtered as described above in order to provide meaningful comparisons to the filtered experimental data.

Table 8. Comparison of measured and predicted SD scores.

Speed	Variable	Data	Model
96 km/h (60 mi/h)	path error (ft)	0.653	0.674
96 km/h (60 mi/h)	wheel (deg)	0.860	0.904
64 km/h (40 mi/h)	path-error (ft)	0.535	0.522
64 km/h (40 mi/h)	wheel (deg)	0.908	0.877

All predicted SD scores were within 5 percent of the corresponding experimental measures. The model correctly predicted the trend of the path errors scores with nominal forward velocity (about a 30-percent higher score for the higher speed). This relatively close match between model and experiment provides at least a partial validation of the model structure.

Sample time histories of lane position (path) and steering wheel deflection from two test subjects are presented in figures 16 and 17. (These time histories were obtained from data taken many minutes into the on-road session and have been assigned a start time of 0 s to facilitate plotting.) Because the external disturbance function is at best known only statistically, we cannot make meaningful point-to-point comparisons of experimental and predicted time histories. (Such comparisons are possible for the laboratory data because the specific time history of the external input is recorded.) Furthermore, the lack of frequency spectral computations for the experimental data prevents a quantitative comparison in that dimension. Nevertheless, visual comparison of the model-generated time histories presented in figure 15 with these data provide a rough qualitative validation of the behavior of the driver/vehicle model.

(a) Path

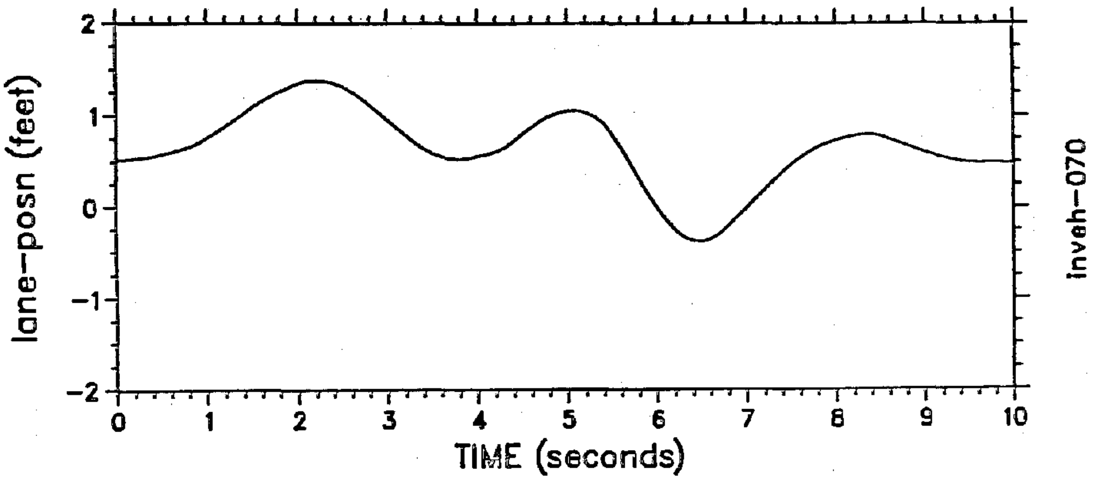

(b) Wheel

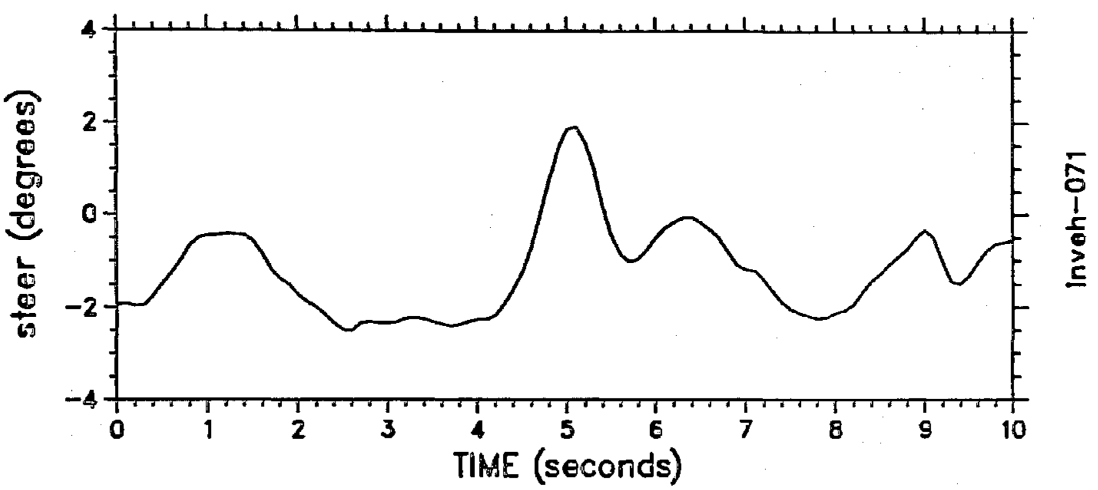

1 ft = 0.305 m

Figure 16. Sample time histories from subject A.

(a) Path

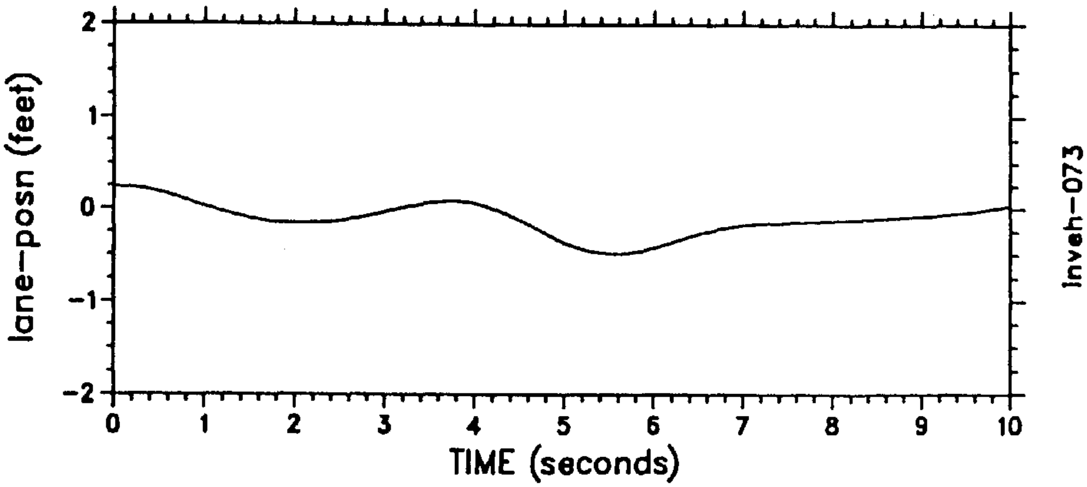

(b) Wheel

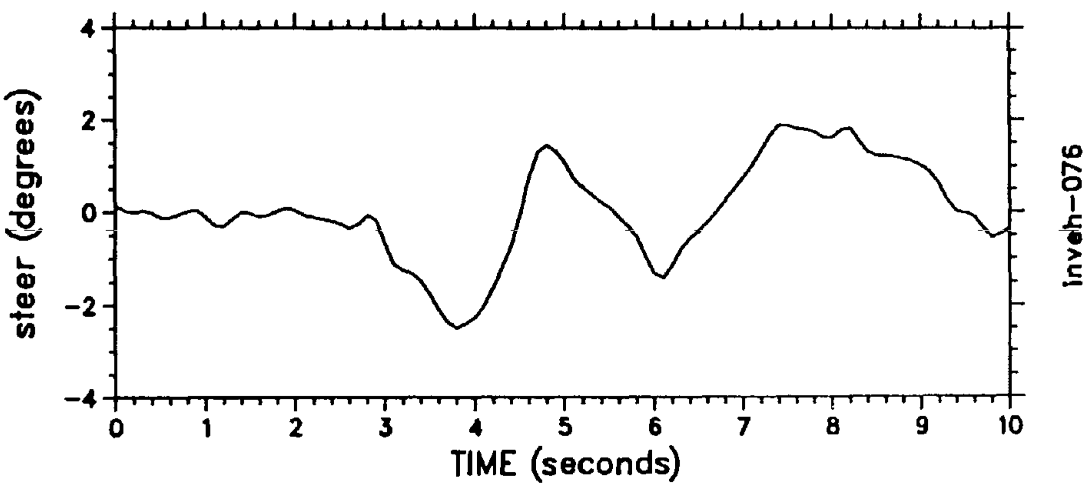

1 ft = 0.305 m

Figure 17. Sample time histories from subject B.

To demonstrate the effects of the wide-bandwidth noise process on the motion of the vehicle, a number of model runs were made in which the vehicle was first initialized with no path, heading, or turn-rate errors, and the noise process was then applied for a simulated interval of 6 s. Steering inputs were omitted in order to guarantee open-loop vehicle response. One hundred time histories were generated each for constant speeds of 64 km/h (40 mi/h) and 96 km/h (60 mi/h).

Ensemble statistics were performed for each set of 100 trials to determine the standard deviation of the lane position as a function of time. The results of this analysis, displayed in figure 18, are consistent with the following approximate theoretical relationship predicted for the models of external disturbance and vehicle dynamics used in this simulation :

$$\sigma_y(T) = KV^2 T^{1.5} \tag{8}$$

where σ_y (T) = standard deviation of the lane position T seconds after application of the noise process, V = velocity in ft/s, and K is a constant proportional to the square root of the power density level of the (approximate) white noise process (Zwahlen and Balasubramanian, 1974).

These results are well within an order of magnitude of the experimental results of Zwahlen and Balasubramanian, who measured the lane deviations of an automobile driven by subjects with their eyes closed following attempts to minimize path and heading errors prior to the period of visual occlusion. Their results were best fit by a mathematical relationship similar to that shown above, except that the velocity term appears to the first power rather than as a squared term. Their mathematical model assumed that the subjects began the period of occlusion with a normally distributed heading error and continued to steer the car during the period of occlusion to minimize perceived lateral and rotational accelerations. In general, Zwahlen and Balasubramanian found

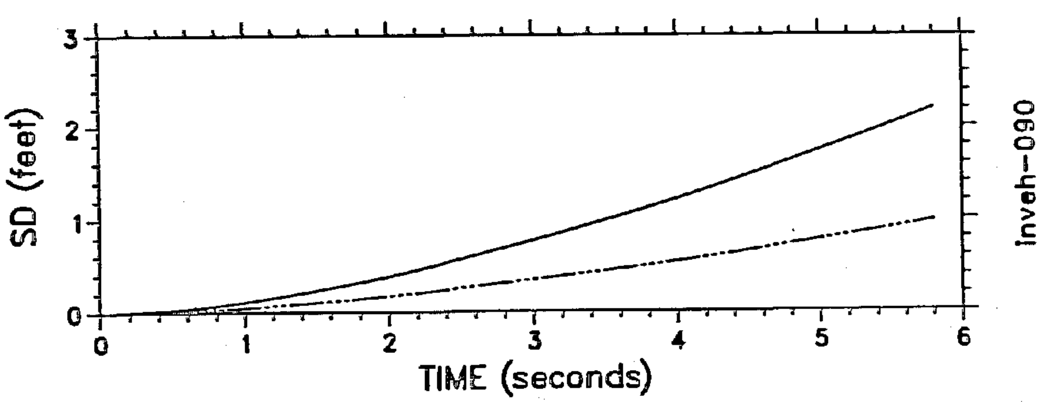

1 ft = 0.305 m

Solid curve = 96 km/h (60 mi/h). Dot-dash curve = 64 km/h (40 mi/h).

Figure 18. Predicted SD of vehicle lane displacement.

larger path errors than shown in figure 18, which is not surprising considering the idealized conditions assumed for the model runs. The degree of correspondence between model and experimental results shown here suggests that the noise level found in the process of calibrating the model is physically reasonable.

VALIDATION AND CALIBRATION OF THE TASK-SHARING MODEL

The foregoing model analysis has focused on calibration of the pilot/vehicle model component in a single-task (i.e., drive-only) situation. This section explores the task- and attention-sharing aspects of the model. Specifically:

- Sensitivity analysis is performed to demonstrate that variations in model parameters preclude qualitatively expected performance trends, and also to support an experimental finding of unexpected performance trends.
- The full capability of the integrated driver model is explored in a model-based simulation of the concurrent driving and telephoning task explored in the UMTRI experimental study.
- Consequences of the assumptions regarding the interaction between attention-sharing and predicted steering performance are explored.

Sensitivity Analysis

Two analyses were performed using the heading-control vehicle model. The first explored predicted model trends as a function of selected task- and driver-related parameters, and the second tested the ability of the model to reproduce some non-intuitive trends reported by Noy (1990). Both analyses employed a simple simulated driving task—lane maintenance on a straight road in the presence of external disturbances (e.g., wind and road surface effects).

Sensitivity to Selected Parameters

The scenario posed for the first analysis was that the driver shared visual attention between the forward scene cues (i.e., the information needed for lane maintenance) and sightseeing (i.e., a non-specific task competing for visual attention). The penalty for not sightseeing was set at a constant value of unity, whereas the criterion level of concern for exceeding the lane boundary was either fixed at 1 percent or varied as an independent variable of the analysis. (A criterion error of 1 percent means that the level of concern for a probability of exceedence of 1 percent equals the level of concern for not sightseeing. See chapter 4 for details concerning the algorithm for task sharing.)

Two parameters of the task-selection model were tested: (1) the prediction interval used in estimating the probability of exceeding the lane boundaries due to inattention to the roadway, and (2) the probability of lane exceedence acceptable to the driver. In addition, the effects of task difficulty were explored by varying the root means square (rms) disturbance level.

The simulated trial time was 60 s. Each model run used the same random number sequence for stochastic variables (external disturbance and operator noise processes) to eliminate performance differences due to factors other than the model parameters of interest.

Predicted scanning behavior and lane-keeping performance are given in table 9. The first data column contains the independent variable of the analysis: prediction interval for data set a, error criterion in terms of acceptable probability of lane exceedence for set b, and rms disturbance amplitude for set c. The remaining columns contain the following model predictions: dwell fraction indicating the fraction of time that the driver attends to the roadway cues; mean dwell time on roadway cues; mean time looking away from the roadway; scan frequency in terms of looks/second at the road; mean path error (i.e., deviation from the lane center); and standard deviation (SD) of the path error computed for the simulated 60-s trial.

The non-zero mean path errors shown in table 9 reflect the effects of finite random sequences. Because the random number generator is theoretically a zero-mean process, predicted lane error should asymptotically be a zero-mean process. In all cases, the predicted means are small compared to the predicted SD's; thus, predicted rms errors are nearly identical to predicted error SD's.

The effects of model parameters on performance trends were as expected. Increasing the prediction interval yielded increased mean dwell times and decreased mean look-away times, and thus increased dwell fractions. The scan frequency trend was less consistent, but tended to increase with prediction interval. Presumably because of the increased dwell fraction, path error decreased with increasing prediction interval.

The reverse trend was found with the error criterion—as the criterion was relaxed, the dwell fraction decreased and lane-keeping performance generally degraded. Because task difficulty was held constant for these two exercises, manipulation of the task-selection model parameters influenced the trade-off between attending to the driving cues and driving performance.

Manipulation of task difficulty, however, did not allow such a tradeoff. Table 9(c) shows that as task difficulty was increased, the driver performed less well while attending more to the roadway cues. This analytic result is qualitatively consistent with data reported in the literature (Noy, 1990).

The scanning behavior was relatively consistent: the scanning rate generally increased with increasing task difficulty or more restrictive subjective performance requirements.

Summary of Noy's Results

Noy (1990) reports the results of a series of experiments exploring multi-task behavior in a driving simulator. His results are particularly useful for providing a qualitative test of the integrated driver model in that they include measures of task performance as well as a usable set of eye-movement statistics.

Noy's subjects were required to maintain lane position and headway while driving a simulated truck along a roadway having randomly sequenced straight and curved segments. In some of the experiments, the subjects were required to perform a concurrent visual search task in which the subjects looked for a short line in a set of otherwise uniform long lines. Experimental variables included: (1) difficulty of the driving task as determined by the curvature of the road segments; (2) difficulty of the search task as determined by the number of lines presented; and (3) relative importance of the auxiliary task, which the experimenters attempted to manipulate via instructions to the subjects. Driving was continuous, and the search task was self-paced; i.e., once the subject indicated presence or absence of a deviant line, a new screen was presented.

Table 9. Parametric study of the decision algorithm.

Independent Variable	Mean Dwell Fraction	Mean Dwell Time (s)	Mean Look-Away Time (s)	Scan Freq (1/s)	Path Error: Mean (ft)	Path Error: SD (ft)
(a) Effects of Prediction Interval (seconds)						
1.0	0.19	0.33	1.51	0.56	0.21	1.34
1.4	0.34	0.41	0.81	0.83	0.18	0.86
2.0	0.84	1.07	0.20	0.79	0.11	0.66
(b) Effects of Error Criterion (allowable probability of lane exceedance)						
1%	0.34	0.41	0.81	0.83	0.18	0.86
2%	0.25	0.31	0.97	0.79	0.09	1.18
5%	0.19	0.33	1.44	0.58	0.21	1.10
(c) Effects of rms disturbance amplitude (feet)						
5.0	0.13	0.31	2.07	0.43	0.24	0.85
10.0	0.34	0.41	0.81	0.83	0.18	0.86
20.0	0.84	0.67	0.21	1.14	0.19	1.15

1 ft = 0.305 m

Driving performance measures included standard deviation of the lane position, time-to-line crossing (a measure derived from vehicle position, velocity, and relative turn rate), and headway. Auxiliary task performance was defined in terms of the time required to respond following the onset of the presentation. Typically, subjects required a number of scans to complete the search task.

Principal findings of the Noy study included:

- Driving errors were larger under dual-task than under single-task conditions.
- Manipulating the difficulty of the driving task had a modest effect on driving performance, a larger effect on scan behavior, but little effect on secondary task performance.

- Manipulating the difficulty of the auxiliary task had the opposite effect: there were only small effects on driving performance and scanning, but relatively large effects on auxiliary task performance.

- When the concurrent-task driving performance data were aggregated according to where the subject was looking (i.e., at the road or at the auxiliary display), analysis showed that driving performance was, on the average, slightly better when the subjects were looking at the auxiliary display than when looking at the road.

This last result at first appears counter-intuitive. Noy explains it by hypothesizing that since visual attention was under the control of the subjects (rather than forced by the experiment), the subjects would look away from the road when they felt that the vehicle state was satisfactory and would look back at the road when errors were starting to build up. This hypothesis is entirely consistent with the philosophy underlying the decision module of the integrated driver model. Implied (but not tested by Noy) is that the reverse trend would be found if the driver were to adopt a scanning strategy (e.g., periodic scan) that did not consider instantaneous vehicle state.

Because of the detail in which performance measures were reported, Noy's data base provides a unique opportunity to validate, at least qualitatively, the behavior of the integrated driver model.

Model runs were formulated and analyzed to determine whether or not the model would replicate two experimental trends: (1) the degradation in driving performance when a concurrent auxiliary monitoring task is required, and (2) the better performance found when the driver was not looking at the road, given the driver's ability to control his scanning behavior. To provide an additional check on the reasonableness of the model predictions, the model was run under the assumption of a periodic scanning behavior.

The driving task explored in the model was the same as that used in the preceding sensitivity analysis. Analysis allowing the driver to share attention according to the task-sharing algorithm implemented in the IDM was first performed. Dwell times for the periodic-scanning trials were subsequently selected to provide a close match to the average times predicted for the driver-controlled scanning condition. The simulation update rate for the task-selection model was 0.2 s.

Because the driver/vehicle model uses random noise sequences to generate the external disturbance process as well as various sources of driver response randomness, model predictions obtained in a single trial are somewhat dependent on the specific random number sequences. In order to assess the reliability of any performance trends that might be seen, four 60-s trials were run with different random sequences, and means and standard errors for each performance variable of interest were computed.

Average scanning behavior predicted under conditions of driver-controlled attention-sharing, along with parameters imposed for the periodic-scanning trials, are shown in table 10. Values enclosed in parentheses are estimated standard errors of the mean.

Average predicted path error standard deviation (SD) scores are shown in table 11 for the three attention-sharing conditions explored with the model. Because predicted mean errors were relatively small compared to the SD scores, the SD scores may also be considered as predictions of rms path error. Three scores are shown for multi-task conditions: the overall score computed from the full data set, the on-road score computed from error predictions corresponding to times when the driver is predicted to be looking at the road, and the off-road score computed from data correlated with attention away from the road.

Table 10. Average predicted scanning behavior.

Attention-Sharing Mode	Scan Frequency (looks/s)	Dwell Fraction	Dwell Time (s)	Away Time (s)
Controlled	0.87 (0.03)	0.32 (0.01)	0.37 (0.02)	0.78 (0.02)
Periodic	0.83 (0.00)	0.33 (0.00)	0.40 (0.00)	0.80 (0.00)

Table 11. Average path error SD (ft).

Attention-Sharing Mode	All Data	On-Road	Off-Road
Controlled	1.03 (0.07)	1.20 (0.11)	0.94 (0.06)
Periodic	0.98 (0.06)	0.96 (0.06)	0.99 (0.06)
Full Attention	0.67 (0.04)	---- ----	

1 ft = 0.305 m

Table 11 shows the following predicted trends: (1) driving performance is worse when attention is shared between roadway cues and some unrelated monitoring task, compared to driving alone; (2) when attention is controlled by the driver, performance is better, on the average, when the driver is looking away from the road than when attending to driving-related cues; and (3) when scanning is forced to be periodic, performance is slightly better when attending to the road. The mean and standard error scores shown in this table suggest that the first two trends, but not the third, would have been statistically significant had these been experimental data.

The model results previously presented, combined with the results reported here, yield the comparisons of predicted and experimental (by Noy) performance trends shown in table 12. Trends are shown for driving performance and scanning behavior. No trends are shown for auxiliary task performance, as this was not a factor in the model analysis.

Table 12 shows that where model and experimental results are available for comparison, the model mimicked all but one of the experimental trends. This discrepancy, along with other aspects of the model results, are discussed below.

The major discrepancy between model analysis and experiment concerns the effects of manipulating the relative importance of the auxiliary task. The model yielded the expected result: attention to the driving task decreased and driving performance degraded as the relative importance of the auxiliary task increased. Noy's experiment, however, showed no significant effect on scanning behavior or driving performance. To some extent, this result may have been due to experimental technique because drivers were generally encouraged to maintain good driving performance in the presence of a side task. Noy suggests, however, that drivers have learned how much attention is generally required by the driving task in a given situation, and that they are reluctant to reduce their attention below the required amount. This assumption can be accounted for in the IDM by assuming a fixed penalty function for the non-driving task, independent of the specifics of the auxiliary task.

Table 12. Predicted and measured performance trends.

	Performance Trends	
	Model	**Experiment**
Multi-task vs single-task driving performance	Multi-task worse	Multi-task worse
Increasing the difficulty of the driving task	More attention to driving task Driving performance degraded	More attention to driving task Driving performance degraded
Increasing the relative importance of the auxiliary task	Less attention to driving task Driving performance degraded	No significant influence on attention or driving performance
Driver-controlled attention	Better performance when not attending to road	Better performance when not attending to road
Periodic attention	Better performance when attending to road	(no data)

Although experimental trends were qualitatively replicated, the magnitudes of on-road and off-road performance differences were somewhat surprising. Specifically, the model predicted that the on-road error scores would be about 25 percent greater than the off-road scores, whereas Noy appears to have found a smaller difference. Similarly, although the model predicted the reverse trend for periodic scanning, the differences were very small. In general, one would expect performance to be substantially worse for periods of inattention to the road when attention-switching is not optimized by the driver. Nevertheless, as shown later, the most deleterious effects of momentary inattention to the road are predicted to occur a short time after attention has been re-directed to the road, even when scanning occurs independent of the instantaneous state of the system. This phenomenon may have contributed to the predicted trends shown in table 11.

To some extent, the model predictions may have been biased by the relatively fast scan frequencies allowed. Table 10 shows a scan frequency of nearly 0.9 looks/s for the driving task environment assumed in this analysis. Noy's data, on the other hand, showed an average scan frequency of around 0.35 looks/s, with most of the data falling between 0.25 and 0.45 looks/s. A slower scan rate would be expected to bias good performance more toward intervals of attention to the road.

The model analysis reported here assumed very little in the way of constraints on scanning rate. The only hard constraint was imposed by the update rate of 0.2 s imposed on the task-selection model. Consequently, some of the predicted dwell times were only 0.2 s long. Two additional constraints were imposed on subsequent model analysis. As indicated in chapter 4, a "dead time" of 200 ms was imposed on eye movements between the on-road and in-vehicle look points, during which time the driver was assumed to obtain no useful information relevant to either the driving or auxiliary tasks. Second, a commit time greater than the base update rate (400 ms) was imposed to maintain some reasonable minimum amount of attention to a given task following the switching of attention from one task to another.

Predicted Effects of the Telephone Task on Driver Performance

The model was used to predict the effects of a concurrent in-car telephoning task on lane-keeping performance and visual scanning behavior. The telephone task implemented in the IDM was intended to replicate as closely as possible some of the experimental conditions explored by UMTRI in their laboratory study. Model analysis was performed using the path-control vehicle model.

The model analysis described here was performed without knowledge of experimental results and may thus be regarded as *a priori* predictions. This exercise, then, is not a true calibration, but rather a demonstration of model application to a multi-task situation. Comparisons are made between model predictions and corresponding experimental data.

A summary of the related laboratory experiments is presented in Serafin, Wen, Paelke, and Green (1993).

Task Analysis

The experimental program explored two levels of each of three experimental factors, yielding a total of eight experimental variations of the telephone task. The three factors were:

- Manual versus voice-operated dialing.
- Local calling (a total of 7 digits to dial) versus long distance (a total of 11 digits).
- Familiar versus unfamiliar telephone number.

In the case of a familiar number, the driver was prompted with the name of the person or place to be called and then dialed from memory. Both name and number were provided when the telephone number was unfamiliar to the driver.

The model was run for the baseline driving-only condition described previously (simulation of laboratory driving task), and for the following three configurations of the telephone task: (1) long-distance (LD), familiar; manual dialing (2) LD, unfamiliar, manual dialing, and (3) LD, familiar, voice dialing. The LD, familiar, manual-dialing task as modeled here consisted of the following steps:

1. Flip to the next page of an index-card telephone directory located on the seat next to the driver.
2. Read the name of the person to be called. (For this condition, the subject knew the telephone number from memory and needed to be prompted only as to the person to be called.)
3. Pick up the handset.
4. Dial 1 + area code.
5. Dial the exchange.
6. Dial the last four digits.
7. Press the "Call" button.
8. Conduct a 30-s conversation.
9. Press the "End" button.
10. Set the telephone down.

Model analysis was based on the assumed performance times shown in table 13 for the component activities underlying the 10 steps listed above. A hand (thumb) movement of at least 200 ms was assumed prior to entering each digit of the telephone number, and an additional 200 ms was assumed for entering the digit. Minimum dialing times were thus 1.2 s for the three-digit exchange and 1.6 s for four-digit combinations.

Except for the 30-s conversation, the (mathematical) driver was required to share visual attention between the simulated roadway scene and a display or control associated with the telephone task.

The driver was assumed to perform each of the visual/manual task segments in the above list without interruption. When a task was completed and a subsequent visual and/or manual task became current, the model determined whether the driver should continue with the next telephone task segment or glance at the road for at least a pre-determined minimum interval (400 ms). Similarly, if the driver was looking at the road while an in-vehicle visual task was pending, the model determined whether or not it was appropriate for the driver to resume the telephone task sequence. (The decision algorithm is described in chapter 3.)

Table 13. Assumed performance times.

Activity	Performance Time
Eye movement	200 ms
Hand movement	200 ms between buttons on the handset
	400 ms between devices
Flip page	1 s
Read name	800 ms
Pick up, set down handset	400 ms
Push button	200 ms

The driver was assumed to obtain roadway cues while conversing, but with some performance decrement. While the conversation was in progress, the model allocated cognitive attention between the driving and conversational tasks according to instantaneous driving performance. (See Levison, 1979, for a discussion of how the effects of cognitive attention-sharing on continuous control performance are modeled.) Attentional allocation was quantized to the nearest 0.25, with the restriction that the driver always allocated at least 25-percent attention to the driving task.

When dialing an unfamiliar number, the driver was assumed to remember only three or four digits at a time. Accordingly, the task analysis for the LD, unfamiliar, manual telephone task differed from the analysis presented above in that the driver was required to read one block of numbers immediately before dialing the 1+area code, another before dialing the exchange, and again before dialing the final four digits.

To model the LD, familiar, voice-operated configuration, the task analysis shown above was modified by eliminating the manual dialing sequence and inserting a voice-dialing sequence in which the driver was assumed to take 1 s to speak each digit. As was the case with the conversation, the driver was assumed to be looking at the road while speaking the digits and to be sharing cognitive attention between the voice-dialing and steering tasks.

Four model runs were obtained for each of the 4 conditions explored in this model analysis—the baseline (drive-only) condition plus the 3 experimental configurations described above—making a total of 16 runs. Each model run represented 60 s of simulation time and included one exposure to the telephone task starting 5 s into the run. The telephone task was typically completed between 45 and 50 s into the run, with the remaining time devoted entirely to the driving task. (Completion time was a variable that depended on how attention was shared between the driving and telephone tasks.)

Run-to-run variability within a condition was obtained by varying the random sequence used to produce the stochastic portion of the driver's steering response. The same set of random sequences was used for each condition. That is, replication #1 for each of the four conditions used one random sequence, replication #2 for each condition used another sequence, and so on. This technique minimized the masking influence of run-to-run variability on main effects in a manner that cannot be done with live subjects and thereby provided statistical significance with less simulation time than would likely be needed in a manned simulation experiment.

Predicted Lane-Keeping Performance

Figure 19 shows the path error time histories for one of the baseline runs and for the corresponding run obtained for the telephone task. Because the same random number sequence was used to produce the stochastic portion of the driver's response, the differences between the two curves reflect the effects of the concurrent telephone task and are not confounded with run-to-run variability.

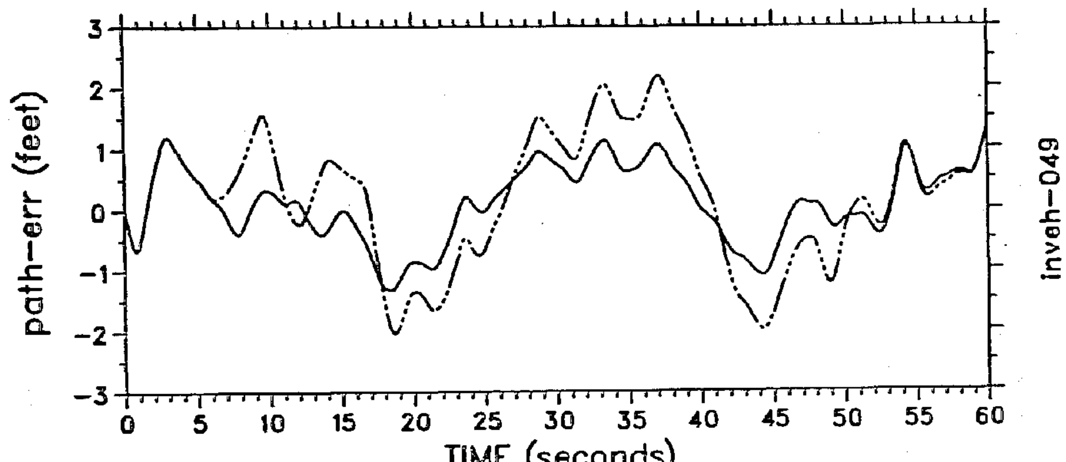

1 ft = 0.305m

Solid Curve = steering only. Dot-dash curve = steering + telephoning.

Figure 19. Predicted path error trajectories for the baseline and experimental conditions.

As expected, larger error peaks are shown for the experimental task. Comparison of this figure with the cognitive attention timeline of figure 20 shows that performance decrements were predicted for the portions of the task when conversation was in progress, as well as when visual attention was shared between driving and telephone start-up tasks. For this curve, a value of 1 indicates full attention to the road, a value of 0 indicates full visual attention to the telephone task, and values between 0 and 1 indicate a reduced level of cognitive attention to the roadway cues while conversing. According to this timeline, the driver made three minimal 0.4-s glances and one 0.8-s glance to the road while performing the pre-conversation telephone tasks.

Within-run standard deviation (SD) scores were computed from the predicted path error and wheel displacement time histories for each model run. To allow for initial transients to subside, the analysis was begun 5 s into the run and continued until the run terminated at a simulated 60 s. (Because the road curvature and random noise processes are theoretically zero-mean processes, the SD score is nearly equivalent to the root-mean-squared error.) Across-trial means and standard deviations of these scores were then computed.

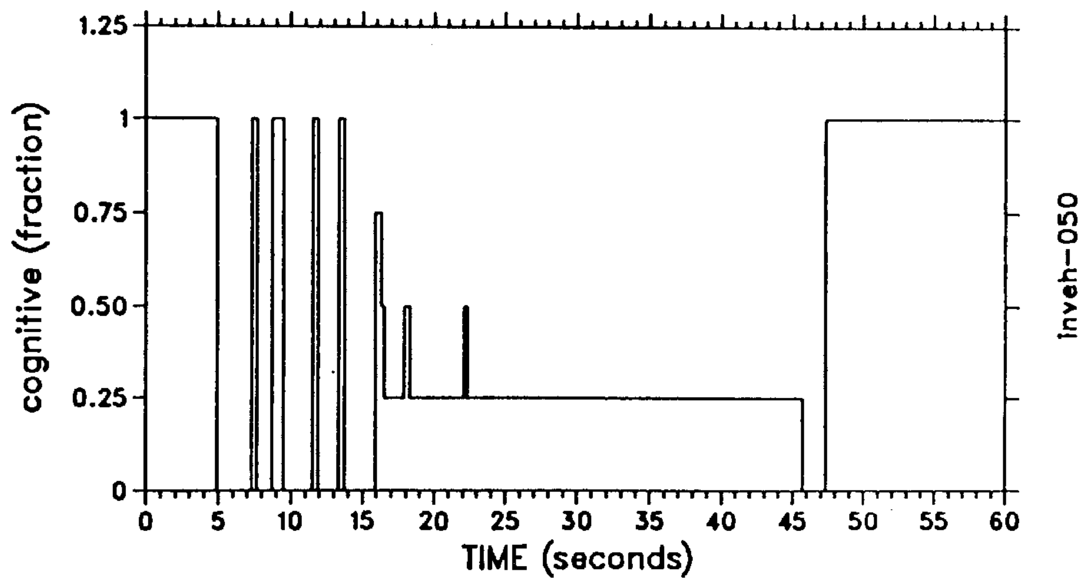

Figure 20. Predicted attention to the driving task while performing the concurrent telephone task.

Table 14 shows that the predicted path error scores for the combined driving and telephoning tasks were about twice the scores predicted for the drive-only baseline condition. T-tests comparing the baseline score with each of the experimental conditions revealed alpha levels of significance of less than 0.01 for each comparison. Differences among the three experimental conditions were not significant (alpha greater than 0.05). Predicted wheel displacement scores for the experimental conditions were about 20 percent greater than the baseline scores and were significant at the 0.01 level. For the particular driving task studied here, the larger path error scores obtained with the concurrent telephone tasks were still relatively small compared to the assumed 1.22-m (4-ft) clearance between the side of the car and the lane boundary for a car centered within the lane.

The predicted results were not entirely as expected. First, there were no significant differences among the performance effects of voice- and manually dialed configurations. Second, error scores were, on the average, somewhat greater for the situations where the number to be dialed was familiar than when it was unfamiliar, even though (as we show later) the model predicted more time spent visually scanning the telephone number pad when the number was unfamiliar.

Table 14. Predicted standard deviation scores.

	Drive Only	LD Fam Manual	UD Unfam Manual	LD Fam Speech
(a) Path Error Standard Deviation (feet)				
Mean	0.65	1.34	1.25	1.19
SD	0.06	0.34	0.22	0.26
(b) Wheel Standard Deviation (degrees)				
Mean	19.0	23.8	23.3	22.8
SD	0.5	1.6	1.1	1.9

Number of Subjects = 4

1 ft = 0.305 m

One might suppose that with the ability to maintain visual attention to the road while voice-dialing the telephone, the voice-operated telephone would have been predicted to be substantially less disruptive of the steering task than the manually operated device. It should be noted, however, that fully half of the model run was devoted to the 30-s conversation, which would be expected to have the same average impact on driving whatever the dialing mechanism. Inspection of the predicted path error timelines shown in Figure 19 shows qualitatively that the degradation in path regulation during the conversation interval (from about 25 through 55 s into the run for this example) was on the order of that observed for the pre-conversation tasks. Furthermore, the model tended to allocate only about 25 percent of the cognitive attention to the road when the driver was concurrently driving and performing a speaking or listening task. The interference between voice-dialing and driving was, therefore, on the order of the interference between manual dialing and driving, where cognitive attention to the road fluctuated between 0 and 1 according to the visual scanning activity.

The allocation of cognitive attention between the driving and auditory tasks is determined, in part, by the settings of the parameters in the penalty function. Should empirical evidence imply that a greater fraction of available attention is paid to the forward scene cues, the penalty weighting coefficients could be revised to reflect this assumption. The effect on predicted driving performance would be to reduce predicted path (lane-keeping) errors for the manually dialed telephoning tasks (by lessening the interference due to the 30-s conversation), and to reduce even more the path errors for the voice-dialed telephoning tasks.

The unanticipated prediction that driving performance would be slightly worse when a familiar number was dialed appears to have been the result of a particularly large sample of observation noise occurring in one of the model runs at an inopportune time when the driver was predicted to be making a minimal (0.4-s) glance to the forward scene. When the model was rerun with the minimum glance time set to 0.6 s, the predicted error SD score for this replication was reduced by 25 percent, virtually removing the predicted (average) difference in error scores for the familiar- and unfamiliar-number conditions.

Predicted Attention-Sharing Behavior

The predicted fractions of attention to the roadway cues are given in table 15. Each entry represents the mean of four such computations, one for each replication of a given condition, computed between 5 and 60 s. The visual attention was defined as the fraction of time that the eyes were predicted to be on the road. (The times in which the eyes were in motion, either to or away from the road, were considered as attention away from the road.) The fraction of cognitive attention was computed by accumulating the fractional allocations of attention for each 0.2-s interval in which the driver was looking at the road and dividing by the number of intervals in the 55-s computation period.

Table 15. Fraction of attention allocated to roadway cues.

	LD Familiar Manual	LD Unfamiliar Manual	LD Familiar Voice
Visual	0.81	0.78	0.92
Cognitive	0.41	0.39	0.36

The trend of the predicted allocation of visual attention is what we might expect intuitively. The greatest amount of visual attention to the road is predicted for the voice-dialed telephone, the next highest for the manually dialed telephone when the called number is familiar to the caller, and the least when the driver is unfamiliar with the number and needs to spend additional time reading the telephone number.

As expected, a larger fraction of cognitive attention was devoted to the roadway cues when the telephone number was familiar to the driver than when the number was unfamiliar. Because cognitive attention to the driving task is assumed to be zero when the driver is looking inside the vehicle, it follows that the fraction of cognitive attention to the roadway would be less for the unfamiliar-number case because of the greater fraction of time spent looking at the telephone number pad.

Table 15 shows that the least amount of cognitive attention was paid to the roadway for the voice-dialed configuration. Even though the driver continued to look at the road while voice-dialing (speaking) the telephone number, the model predicted that the split of cognitive attention during concurrent driving and speaking would generally be around 25 percent to the roadway cues and 75 percent to speaking or listening.

Predicted scanning statistics are shown in table 16 for the telephone tasks with manual dialing. All variables, except for the total fraction, were computed for the interval containing the subtasks preceding the conversation (tasks 1 through 7 listed above). In the following discussion, all visual targets related to the telephone task (directory, handset, number pad) are referred to simply as the "telephone." The predicted variables are as follows:

- Duration. Interval to complete the pre-conversation tasks, in seconds.
- No. Scans. Number of scans (glances) at the telephone generated while performing these tasks.
- Scan Frequency. Number of glances/second at the telephone.
- Total Fract. Total time visually attending to the telephone, normalized with respect to the entire scoring interval of 55 s.
- Dial Fract. Total time attending to the telephone, normalized with respect to the duration of the pre-conversational tasks.
- Dwell. Average time in seconds attending to the telephone during a single glance, in seconds.
- Interrupt. Average time attending to the road between successive glances at the telephone, in seconds.

Table 16. Predicted mean scanning behavior for use of manually dialed telephone.

	Duration	No. Scans	Scan Freq.	Total Fract.	Dial Fract.	Dwell	Interrupt
Familiar	12.6	5.25	0.42	0.13	0.65	1.61	1.06
Unfamiliar	14.0	6.25	0.45	0.15	0.65	1.45	0.95

Attention was computed somewhat differently for this analysis than for the results presented in table 15. "Attention" was defined here in terms of the task that the model selected for processing. That is, a "dwell" interval was defined as the interval between the time the model determined that the driver should attend to the telephone task until the next time that the model determined the driver should attend to the roadway cues. An "interrupt" interval was define in a similar manner. Thus, attention to a task included any preparatory times to move the eyes and/or free hand (which is typically the case when analyzing actual eye-movement data).

All entries shown in table 16 represent means computed across the four replications for the experimental condition. One measurement of each of the variables shown in the first five columns was obtained per model run. To compute mean dwell and interrupt times, however, within-run means were first computed, and these means were then averaged across replications.

As expected, more scans and more total time were predicted to complete the pre-conversation tasks when the telephone number was unfamiliar to the driver. Consequently, the total fraction of time spent looking at the telephone was greater when dialing an unfamiliar number. While the pre-conversation tasks were in progress, however, the fraction of attention diverted from the roadway was about 65 percent for the two conditions, suggesting an equal disruption of the driving task. Predicted scanning frequency was somewhat greater for the unfamiliar-number condition, with reductions in both average dwell and interrupt times.

Comparison of Model Predictions with Experimental Data

Data provided by the laboratory experiments are compared with model predictions in table 17.

Table 17. Comparison of predicted and experimental performance measures.

	Path Error SD (ft)		No. Scans		Mean Dwell IP (s)		Mean Dwell Road (s)	
	Model	Expt.	Model	Expt.	Model	Expt	Model	Expt
Drive Only	0.65		--	--	--	--	--	--
LD Manual Fam	1.34		5.25	6.0	1.61	0.84	1.06	0.94
LD Manual Unf	1.25		6.25	7.8	1.45	1.04	0.95	1.17
LD Voice Fam	1.19		--	--	--	--	--	--

Fam = familiar number, Unf = unfamiliar number

1 ft = 0.305 m

Attention-Sharing Aspects of the Driver/Vehicle Model

The preceding material has focused, in part, on the switching of attention between tasks and on the consequences of attention-sharing on steering performance. The following discussion explores in more detail the assumptions within the driver/vehicle model component as to how attention-sharing affects control performance. In this section, the phrase "attention model" refers to this aspect of the driver/vehicle model element—not to the decision algorithms that determine which task the driver attends to at any instant.

The baseline IDM assumes that when the driver diverts attention away from the roadway, he/she continues to execute steering movements for a short time following occlusion in order to minimize the errors most recently observed. (See appendix A for a mathematical description of the state estimation process.) As the driver continues to be without visual cues, the perceptual information becomes outdated, uncertainty grows, and control movements effectively cease until the forward scene cues are again observed. Stated another way, the driver is assumed to formulate and generate an "open-loop" movement program that is designed to reduce path error, based on the information available at the time of occlusion.

One may postulate reasonable alternatives to this model. As noted above, the model proposed by Zwahlen and Balasubramanian assumes that the driver maintains control throughout the period of occlusion to the extent of attempting to minimize lateral and rotational accelerations. Another hypothesis is that when the driver is presented with a specific visual side-task, he/she becomes so absorbed in that task that no attention is paid to the driving task and the wheel is held in a relatively fixed position until the forward scene is again observed. In the following development we test this latter "no-steering" alternative against the baseline model. (Resources did not permit a test of the "minimize-acceleration" hypothesis.)

Early experimental results reported by Senders et al. (1967) on driving under conditions of frequent and repetitive intervals of visual occlusion strongly indicate that the driver executes effective steering actions during part or all of the visual occlusion interval. This conclusion is supported by the analysis with the IDM. One of the tasks performed by the drivers in the on-road study of Senders et al. was to drive as fast as perceived safety would allow when vision was alternately occluded for a fixed interval and unobstructed for another fixed interval. When they systematically varied the "look interval" from one set of trials to the next, they found that driver behavior in this experiment was relatively constant for look intervals greater than 0.5 s, and not much different for intervals as low as 0.25 s.

On the basis of these results, the IDM was operated in a forced-scanning mode in which the intervals of attention (to the driving cues) and inattention were fixed, i.e., not subject to the task-selection rules described earlier in this report. The period of inattention was fixed at 2.0 s (one of the intervals explored in the earlier experiments), and the period of attention was varied from 0.1 to 1.0 s. Forward speed was assumed to be 64 km/h (40 mi/h) (all experimental subjects were able to handle a 2-s occlusion at that speed). The heading-control vehicle model was used.

The two competing treatments of attention were explored. The no-steering model was unable to maintain effective control of the vehicle for occlusion intervals of 1.0 s or less. The model thus confirms what we intuitively expect, namely that the operator must be able to execute effective steering action during some portion of the occlusion interval when look times are very short.

A test of the baseline model for the effects of attention-sharing, on the other hand, is consistent with the experimental findings that look times greater than 0.5 s had little effect on driver behavior. A 3-min simulated trial of alternating attention and inattention was performed with the model. Table 18 shows relatively little change in predicted path error SD score as the look time increases form 0.6 to 1.0 s, whereas the score increases substantially for very small look intervals.

Table 18. Effects of look interval on predicted path error SD.

Look Time (s)	Predicted Path Error SD (ft)
0.1	1.80
0.2	1.32
0.4	1.03
0.6	0.87
1.0	0.82

1 ft = 0.305 m

While the baseline attention model is consistent with the early occlusion studies, the question still remains as to whether or not drivers will execute effective steering while attention is diverted to another task that imposes both visual and cognitive demands. In the absence of available data directly bearing on this question, one can use the model to explore the likely consequences of various hypotheses of attention-sharing on predicted lane-keeping performance.

Figure 21 shows predicted ensemble statistics (path-error SD scores) versus time following occlusion for the baseline and "fixed-wheel" attention model, for simulations corresponding to speeds of 96 km/h (60 mi/h) and 64 km/h (40 mi/h). Also shown for reference are the idealized vehicle-excursion statistics originally presented in figure 18.

The driving task was modeled as described earlier in this chapter in the discussion of calibration against baseline data. Each model run consisted of 100 segments in which the driver was assumed to observe the forward scene for the first 10 s and not to view the forward scene during the following 6-s interval of occlusion. The data were analyzed in the same manner as the earlier idealized vehicle response was analyzed; ensemble SD statistics were computed for each 0.2-s recorded interval following occlusion.

Figure 21 shows a significant departure between the predicted vehicle excursions for the two attention models starting at around 1.4 s into occlusion for the 96-km/h (60-mi/h) forward-speed case and at around 2.0 s for the lower-speed condition. Model predictions made with this model for attention-sharing effects will therefore likely constrain glances away from the road to under 2 s, which is consistent with both the results of the Senders et al. (1967) study that shows that occlusion times for some drivers were about 2 s or less at speeds at or above 64 km/h (40 mi/h), and other published results that show that glance times are generally under 2 s (Dingus, Antin, Hulse, and Wierwille, 1988). On the other hand, model analysis has implied that the drivers in the Senders et al. study must have exercised some effective control during the initial period of occlusion.

Additional, highly detailed data would be required to resolve the issue as to the exact nature of control behavior during the period of occlusion. It may be that individual drivers vary in their ability to execute effective control during periods of occlusion.

The curves corresponding to the baseline attention model tend to roughly parallel the idealized-vehicle response at an offset of approximately the baseline SD error performance. This result suggests (if one accepts the baseline attention model) that one might be able, in general, to combine the idealized vehicle results with baseline steering performance and with some assumed scan pattern to predict the effects of visual scanning on steering performance. This assumption is not warranted, however.

Although the curves of figure 21 may give an idea of the relative effects of visual inattention, they are not sufficient for predicting effects of specific scan patterns. First, they do not show the entire effects of even an infrequently occurring single glance away from the roadway. For example, these curves seem to imply that occlusions of less than 1 s should have no adverse impact on driving performance. This is not a correct interpretation. What the curves do indicate is that there may not be a measurable performance degradation within 1 s of the onset of occlusion. What they do not show is the performance degradation that occurs *after* attention has been redirected to the forward scene. This phenomenon is discussed further in the next section.

Second, performance will generally depend on both the particular scan pattern used to obtain information from a given display and the frequency with which such scans are repeated. As Dingus, Antin, Hulse, and Wierwille (1988) have reported, multiple scans are often required to obtain information sources other than the forward scene. To use the model to predict the effects of a specific scan pattern on steering performance, it is therefore recommended that the scan pattern of interest be explicitly included in the model analysis.

The model runs providing the data presented above were generated according to a forced-scan assumption; that is, a 10-s interval of attention to the road was simulated, and then visual cues were removed. The task-selection philosophy described in previous chapters, however, assumes that the driver will divert attention from the forward scene cues only when the urgency of the auxiliary task exceeds that of attending to the driving task. To determine what effect this model implementation would have on the results presented above, the model analysis was repeated using the controlled-scan assumption, which was mechanized as follows:

1. Attend to roadway cues for 10 s.
2. Continue to attend to roadway cues until the penalty for not driving (described in chapter 3) yields an effective prediction of lane exceedence of less than 1 percent (using a prediction time of 1.4 s).
3. Drive without visual cues for an additional 6 s.

Figure 22 compares ensemble SD statistics for time following occlusion for the two scan assumptions for the 94-km/h (60-mi/h) speed condition. Results are shown for both the baseline (figure 22(a)) and fixed-wheel (figure 22(b)) assumptions. The differences between the forced-scan and controlled-scan assumptions are similar for the baseline and fixed-wheel attention models. In both cases, the controlled-scan treatment yields a slightly smaller error score than the forced-scan model for the first second or so of the occlusion interval, and higher errors thereafter.

The initially lower errors predicted for the controlled scan are to be expected, because the attention-switching algorithm (step 2 above) forces the occlusion period to start when the error profile is more benign than average. The explanation for the greater errors produced by this strategy later in the occlusion period is less obvious. It is possible that this behavior is reflecting a law of averages regarding the short-term error statistics. That is, if occlusion is forced to start when the driving task is momentarily easier, then it must be the case (given the statistical characteristics that have been assumed for the driving task) that at some time afterward, the driving task will be harder than average.

The analysis presented below to demonstrate model application explored the effects of specific scan patterns on predicted steering performance. Because of the apparently small differences between the model predictions generated by the forced- and controlled-scan assumptions for realistic occlusion times of 2 s or less, the controlled-scan model was used for this analysis to simplify analysis of model-generated time histories.

(a) **Speed = 96 km/h (60 mi/h)**

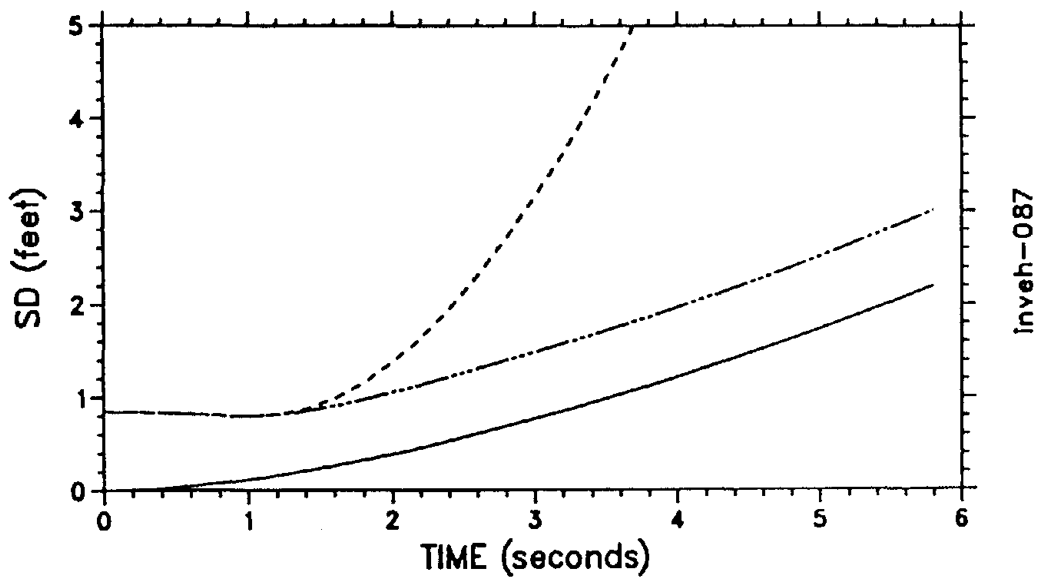

(b) **Speed = 64 km/h (40 mi/h)**

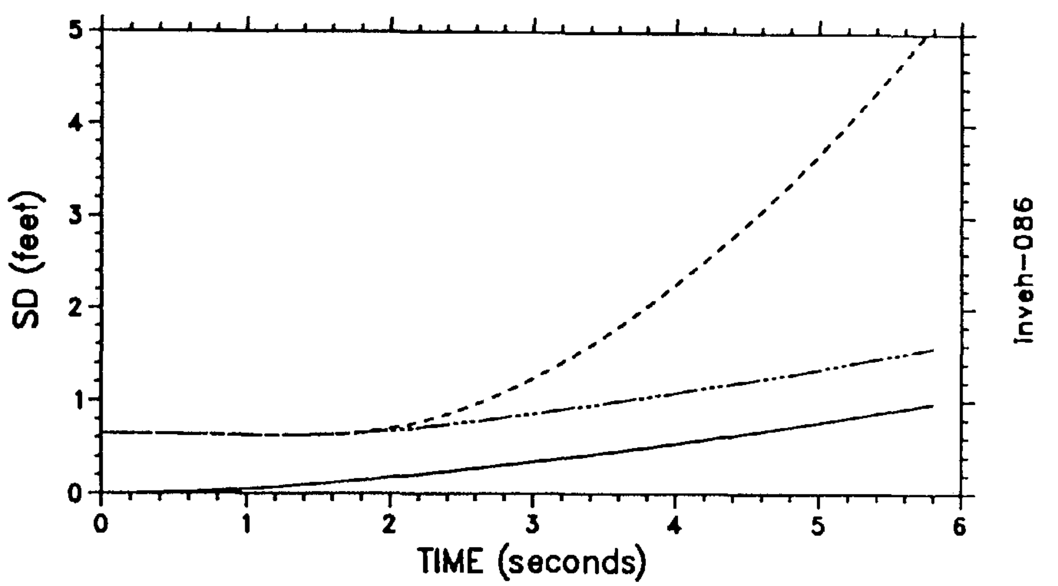

1 ft = 0.305 m

Solid curve: idealized vehicle response, dot-dash curve: baseline attention model, dashed curve: "fixed-wheel" attention model. N=100.

Figure 21. Effects of attention model on predicted vehicle excursions following occlusion.

(a) **Baseline Attention Model**

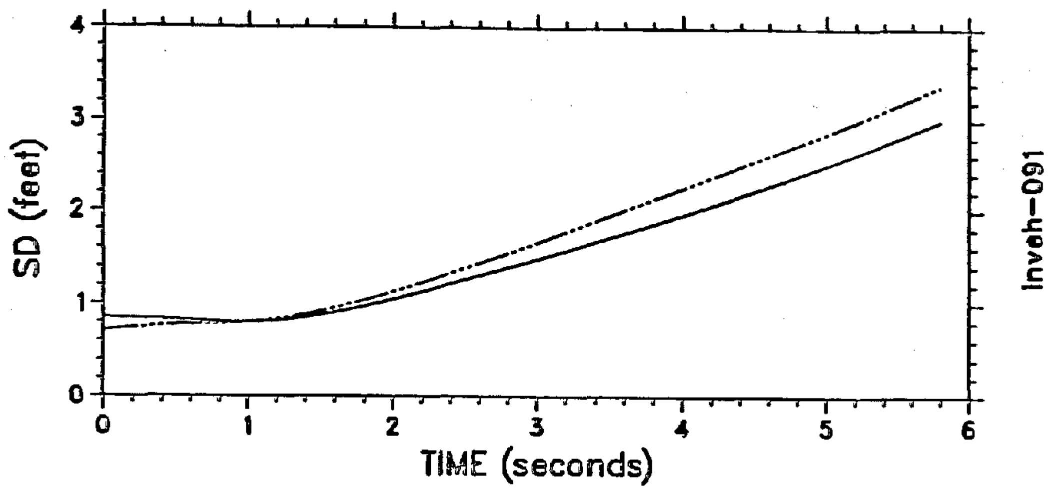

(b) **Fixed-Wheel Attention Model**

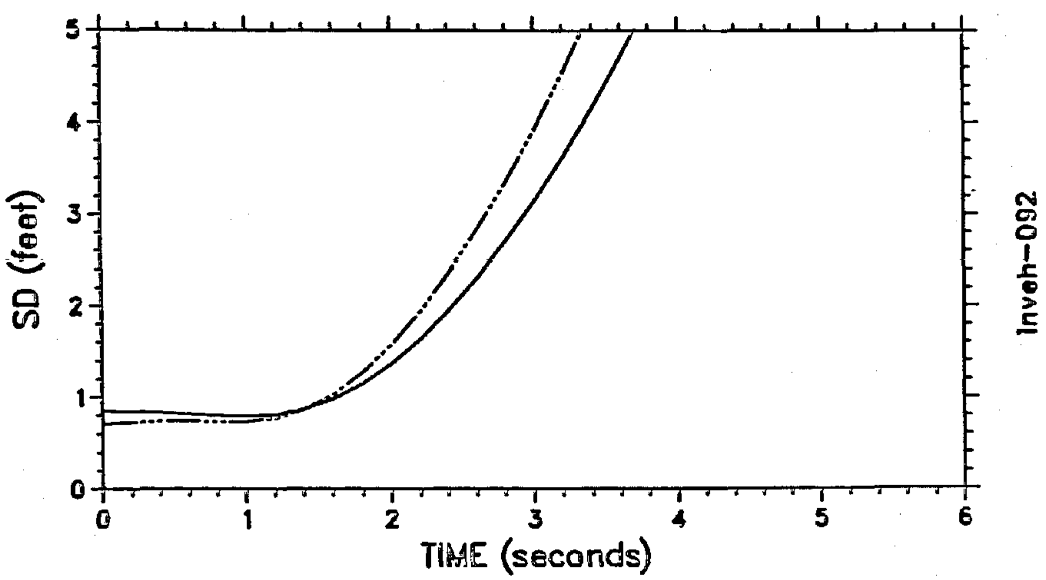

1 ft = 0.305 m

Solid curve = forced scan, dashed curve = controlled scan. N = 100

Figure 22. Comparison of forced-scan and controlled-scan assumptions.

MODEL APPLICATIONS

The foregoing model analysis has concentrated largely on interpretation and support of existing experimental data, model calibration, and the consequences of potential alternatives to the attention-sharing assumptions imbedded in the model. This section demonstrates how the model might be used to generate guidelines for in-vehicle displays. All model analyses summarized below have been obtained using the heading-control formulation for vehicle dynamics.

Based in part on the early work of Senders et al. (1967), Zwahlen, Adams, and DeBald (1988) have proposed guidelines as to what sort of visual scanning requirements may be imposed on the driver by displays and information sources that require diverting attention from the forward scene. Specifically, they suggest acceptable, "gray," and unacceptable combinations of number of glances and time-per-glance for a given glance sequence. These guidelines have been summarized in a companion document to this report (Green, 1993).

In principle, the IDM can be used to generate guidelines of this type and to explore other dimensions of the situation as well (for example, the amount of time intervening between successive glance sequences or differences between visual and auditory displays). Examples of this type of model use are provided in this section. Comparisons of predicted steering performance are made for different visual glance sequences and for different sensory modalities.

The driving task was modeled as described previously for the 96-km/h (60-mi/h) forward-speed condition. Forced-scan timelines (glance sequences) were applied to predict the effects on performance of specific visual scan patterns. Each analysis consisted of a total of 100 exposures to the scan pattern, generated by conducting 2 model runs of 50 exposures each or 4 model runs of 25 exposures each, depending on the length of each exposure.

Two visual glance sequences were explored. The more intrusive sequence consisted of four 2-s glances to an assumed in-vehicle display separated by 0.5-s glances back to the forward scene. This sequence was based in part on the work of Antin et al. (1988) who found a relatively large number of glances per glance sequence to their in-vehicle navigation aids. The display glance time of 2.0 s was within the range of single-glance times reported, and the 0.5-s road glance time was based on the results of Senders et al. (1967), supported by the model analysis reported above.

The less intrusive glance sequence consisted of two 1-s glances to the in-vehicle display separated by a 0.5-s glance to the roadway. This pattern is consistent with the scan data reported by Green, Hoekstra, Williams, Wen, and George (1993) in their laboratory study of in-vehicle navigation aids. It is also consistent with the model suggested in chapter 3 for reading text from a display (i.e., a reading rate of three words/second) if we assume the navigation display contains a six-word message.

A given exposure consisted of an initial period of attention to the forward scene to allow the driving task to approach a statistically steady-state baseline performance condition, followed by the forced-glance sequence. The total exposure time was either 10 s (for the short sequence) or 20 s (for the longer sequence). For example, the timeline for the two-glance pattern was:

- 7.5 s of attention to the forward scene.
- 1.0 s of attention to the in-vehicle display.

- 0.5 s of attention to the forward scene.
- 1.0 s of attention to the in-vehicle display.

Auditory tasks were modeled by constructing an exposure that consisted of an initial interval of full cognitive attention to the forward scene, followed by an interval of reduced attention to the forward scene to represent concurrent processing of visual (roadway) and auditory (e.g., navigation) cues. (See chapter 2 for a discussion of the treatment of cognitive attention within the driver/vehicle model component, and chapter 3 for the treatment of concurrent multi-modality information processing.) The level of reduced attention to the forward scene was a parameter of the analysis.

The auditory analog of the four-scan visual-glance sequence consisted of an initial 11.5 s of full attention followed by 9.5 s of reduced attention. The auditory analog of the two-scan visual-glance sequence was an initial 7.5-s interval of full attention followed by 2.5 s of reduced attention. The period of reduced attention is consistent with an auditory navigation aid that generates a 6-word message at a rate of 150 words/minute. To predict the effects on steering performance when the driver shares cognitive attention between roadway cues and an auditory information source, one must either specify or predict the way attention is allocated to these competing information sources. Algorithms for making such predictions are described in preceding chapters. To use the model in this type of predictive mode, however, one needs to know the tradeoff between attention and performance on the auditory task to construct the appropriate penalty function for use by the task-selection algorithm. In the absence of quantitative data that would allow calibration of this aspect of the model, attentional allocation is treated as an independent parameter of the analysis and, for simplicity, is held constant during intervals when the auditory task is present.

Two model runs of fifty consecutive exposures each were obtained for the conditions having 10-s exposures, making a total of 100 replications. Because the two model runs used different initial seeds for the random noise generation (simulating driver randomness and the stochastic road noise process), every replication represented a statistically independent sample of the driving task in terms of these noise processes. The initial segment of each model run was treated as a warm-up period and was discarded from the analysis, leaving 98 exposures for statistical analysis.

Four model runs of twenty-five consecutive exposures were obtained for the conditions having 20-s exposures, again making a total of 100 replications. Because the initial exposure in each model run was discarded, 96 exposures were used for statistical analysis for these conditions.

Two variations of path-error SD scores were computed for each condition. A global SD score was computed by aggregating all the measurements for the entire 96 or 98 sequences and treating the collection of data points as if they came from one very long replication. Although the within-trial SD score is a standard performance metric for error-nulling manual control tasks, it is not entirely appropriate for the situations explored in this model analysis. In order for such a metric to have the greatest analytical usefulness, the underlying process should be a statistical steady-state Gaussian random process.

The Gaussian assumption is satisfied by the way in which the driving task has been modeled. The statistics of the Gaussian process should not be assumed to be stationary, however, because of the differences in the amount of attention devoted to the forward scene as each exposure unfolds. In general, one would expect errors to be statistically greater during or shortly after scanning activity than during relatively long intervals of full attention to the forward scene. As shown below, the global SD score tends to mask the performance differences across conditions, particularly when the period of uninterrupted attention is large compared to the intervals where attention is being shared between driving and non-driving tasks.

A second computation (actually, a series of computations) generated ensemble SD scores for each 0.2-s sample interval throughout the exposure period. This sequence of SD computations was then examined to find the numerically largest value and the time at which this value occurred. These measures are called the maximum SD and time of maximum SD in the following discussion. To provide an indication of the practical importance of this error measure, the predicted standard deviation was interpreted in terms of the probability of exceeding the lane boundary (i.e., a path deviation of 0.91 m (3-ft) or greater). This probability is termed the maximum probability. Note that this metric is not the fraction of time that the vehicle is predicted to exceed the lane boundary over the course of the simulation, but rather the probability of lane exceedence at the specific time of (statistically) greatest performance decrement.

The maximum statistical-error measure does not provide a full accounting of the effects of attention-sharing, however. Because this measure is a "snapshot" of performance statistics at one specific time following initiation of the scan pattern, it addresses the question of how bad the performance is likely to be. It does not directly indicate the time over which performance is likely to be of concern. To address this issue, another metric was computed from the model-generated time histories: the total time that predicted path error exceeded the assumed allowable error over the accumulation of all 96 or 98 exposures. This metric was then normalized by the number of exposures and multiplied by 10 to indicate the time out per 10 exposures.

Predicted performance metrics are shown in table 19 for the following task conditions:

- <u>Baseline</u>. Full attention to the driving task to provide a baseline performance reference for the following conditions.
- <u>1-s scan</u>. Exposures consisting of 9 s of attention to the forward scene and 1 s of attention to an (undefined) in-vehicle display were analyzed to test the implication of figure 21 that an occlusion of this duration would have no measurable influence on steering performance. The baseline (non-fixed wheel) attention model was assumed.
- <u>Two glances, baseline attn</u>. The less intrusive glance sequence described previously with 1-s glance times to the display, 0.5-s glance times to the roadway, and the baseline model for attention.
- <u>Two glances, fixed-wheel</u>. Same as above except for the fixed-wheel attention model.
- <u>Four glances, baseline attn</u>. The more intrusive glance sequence described previously with 2-s glance times to the display, 0.5-s glance times to the road, and the baseline attention model.
- <u>Four glances, fixed-wheel</u>. Same as above except for the fixed-wheel attention model.
- <u>2.5-s aud., 0.5 attn</u>. The auditory analog of the two-glance visual sequence described previously, with an assumed equal allocation of attention to the driving and listening tasks.

- 2.5-s aud., 0.75 attn. Same as above, except for an assumed attentional allocation of 75 percent to the auditory task and 25 percent to the forward scene.
- 9.5-s aud., 0.5 attn. The auditory analog of the four-glance visual sequence described previously, with an assumed equal allocation of attention to the driving and listening tasks.
- 9.5-s aud., 0.75 attn. Same as above, except for an assumed attentional allocation of 75 percent to the auditory task and 25 percent to the forward scene.

A comparison of the baseline and 1-s scan conditions supports the following predictions and inferences made earlier:

- Performance statistics are relatively stationary when attention to the forward scene is constant.
- The global SD score tends to mask performance differences across conditions.
- One should expect some performance decrement accruing from a 1-s diversion from the forward scene.

The difference between the maximum SD score and the global SD score for the baseline condition is about 8 percent, which provides a measure of the reliability of this particular statistic. The global SD score for the scanning condition is also about 8 percent greater than the baseline score, which is clearly not a significant effect. On the other hand, the maximum SD score is about 33 percent greater than the baseline maximum SD score, which in a carefully controlled experiment, would most likely be a statistically significant effect. Table 19 also shows that the predicted time of maximum performance degradation was a little over 1 s after attention was redirected to the forward scene for this condition and for most of the other conditions as well. Therefore, looking only at performance during periods of attention-sharing is not likely to capture the full performance consequences of shared attention.

A comparison of the baseline and fixed-wheel treatments of attention-sharing shows a modest difference for the two-glance sequence (about 35 percent greater maximum SD for the fixed-wheel treatment). On the other hand, the model predicts the vehicle will not be controllable if the four-glance sequence is attempted by a driver whose behavior during inattention to the forward scene conforms to the fixed-wheel model. This model result agrees with the guidelines proposed by Zwahlen, Adams, and DeBald (1988), which show that the combination of four glances and 2-s glance interval is unacceptable.

One should be careful not to overgeneralize this result. As stated earlier, performance will depend on the details of the scan pattern. A less severe impact on performance by the four-glance sequence is anticipated as glance time to the display is reduced and intervening glances to the road are lengthened.

Unlike the model results presented earlier in this chapter for the laboratory tracking situation, the analysis here supports the findings of Serafin, Wen, Paelke, and Green (1993) that auditory displays have the potential to interfere less with driving performance than corresponding visual displays. Tables 20 and 21 compare predicted performance trends for the analogous visual and auditory tasks. Table 20 displays the growth in the maximum SD score as a ratio of the score obtained in the multi-task environment to the maximum SD score predicted for the baseline drive-only condition. Table 21 compares the time out per 10 exposures.

Table 19. Effect of task conditions on predicted performance measures.

Condition	Global SD (ft)	Maximum SD (ft)	Time of Max SD (s)	Maximum Probability (%)	Time Out per 10 Exp. (s)
Baseline	0.72	0.78	4.2	0.0	0.0
1-s scan	0.77	1.04	2.2	0.4	0.0
Two glances, baseline attn.	0.83	1.16	3.5	0.9	0.0
Two glances, fixed-wheel	0.92	1.57	3.7	5.6	0.5
Four glances, baseline attn.	1.06	1.55	8.9	5.4	2.7
Four glances, fixed-wheel	>100	>100	11.5	100.0	97.7
2.5-s aud., 0.5 attn.	0.77	0.93	3.7	0.1	0.0
2.5-s aud., 0.75 attn.	0.87	1.22	4.1	1.4	0.1
9.5-s aud., 0.5 attn.	0.84	1.02	10.7	0.3	0.5
9.5-s aud., 0.75 attn.	0.93	1.20	11.1	1.2	0.6

N = 96 or 98

1 ft = 0.305 m

The two metrics support the hypothesis that the auditory display of a 2- to 5-s message will interfere less with steering performance than the analogous visual presentation (given the assumptions concerning the visual scan pattern) when 50 percent or less of the cognitive attention is allocated to the auditory task. Furthermore, the time-out-of-bounds measure suggests that this

Table 20. Effects of visual and auditory information display on the predicted fractional growth in the maximum SD score.

In-Vehicle Task	Visual display	Auditory display Attn = 0.5	Auditory display Attn = 0.75
2-Glance Sequence	1.5	1.2	1.6
4-Glance Sequence	2.0	1.3	1.5

Table 21. Effects of visual and auditory information display on the predicted time out of bounds per 10 exposures.

In-Vehicle Task	Visual display	Auditory display Attn = 0.5	Auditory display Attn = 0.75
2-Glance Sequence	0.0	0.0	0.1
4-Glance Sequence	2.7	0.5	0.6

side task, whether visual or auditory, should have little impact on steering performance. As the attentional demand of the auditory task increases, however, the predicted performance effects approach that of the visual task.

The same predicted trend is observed for the longer auxiliary task. The impact of the visual side task is substantially greater than for the shorter task (a factor of 2 increase in the maximum SD score compared to baseline and a predicted 2.7 s out of bounds per 10 exposures to the side task). In this example, auditory information presentation is predicted to interfere substantially less than visual presentation even when attention to the side task is as high as 75 percent.

SUMMARY AND DISCUSSION

This section briefly reviews the highlights of the preceding analysis and suggests guidelines for further model application.

Review of the Model Analysis

The driver/vehicle model component was calibrated against driving performance data obtained in laboratory and on-road experiments. Independent model parameters were largely based on either previous model work or engineering judgment and were fixed throughout the calibration process. A parameter associated with the tradeoff between reducing errors and minimizing rapid steering movements served as the independent parameter of the calibration process and was adjusted to provide a best match to measured path error and steering wheel deflection standard deviation

(SD) scores. These scores were matched to within 5 to 10 percent in both cases, and predicted wheel time histories were qualitatively similar to those observed experimentally. Once calibrated for the 96-km/h (60-mi/h) on-road condition, performance scores for the 64-km/h (40-mi/h) condition were predicted to within 5 percent without further adjustment of independent parameters.

As part of the calibration against the on-road experimental data, it was necessary to determine the magnitude of an assumed road-surface disturbance process. The calibrated road-noise model was then used to generate open-loop vehicle response, which was found to be generally consistent with the results of a similar on-road study performed much earlier.

The full integrated driver model (IDM) was subjected to a sensitivity analysis and, where data were available, predicted performance trends were compared with experimentally observed trends. The following experimental trends were correctly predicted by the model:

- Compared to single-task driving, steering performance degrades when an auxiliary task is imposed.
- Increasing the difficulty of the driving task in a multi-task environment results in more attention to the driving task and worse steering performance.
- Contrary to expectations, when attention-sharing between driving and an auxiliary task is relatively frequent, steering performance is slightly better during intervals when attention is paid to the auxiliary task than to the driving task.

On the other hand, the model predicted that increasing (via instructions) the relative importance of the auxiliary task would result in less attention to the driving task and a consequent degradation in performance. The corresponding experiment showed no such trends in either attention or performance. Should this phenomenon be observed for a substantial class of auxiliary tasks, it can be modeled by assigning a fixed penalty to all such auxiliary tasks.

The model was used to predict the effects of a concurrent in-car telephoning task on lane-keeping performance and visual scanning behavior. The telephoning task required obtaining dialing instructions, looking up the number if unfamiliar to the subject, dialing manually or by voice, and conducting a 30-s conversation. The model predicted that adding the telephoning task would increase the path-error SD scores significantly beyond the scores associated with driving only. The difference between predicted scores for the manually dialed and voice-dialed conditions was unexpectedly low, but a reduction in the assumed attentional demand of the listening and speaking tasks would be expected to show greater performance differences between the two dialing modes. The model predicted that the voice-dialed task required less visual attention overall than the manually dialed task, and that dialing a familiar number required less visual attention and fewer scans inside the car than dialing an unfamiliar number.

The consequences of two algorithms for modeling the effects of attention-sharing on steering performance were explored: (1) the baseline attention model in which the driver is assumed to continue effective steering for a short while following diversion of attention away from the forward scene, and (2) the fixed-wheel model in which the driver is assumed to cease making wheel movements during intervals of inattention. Sensitivity analysis showed a substantial difference in predicted performance decrements. Existing data seem to support the baseline model, but one cannot rule out the possibility that some drivers under high workload conditions may adopt a fixed-wheel behavior. Until further research is able to resolve this issue, one course of action is to retain both models to allow predictions of best-case and worst-case scenarios.

A small sensitivity study was performed to compare the consequences of the controlled-scan hypothesis of attention-sharing where attention is switched from one task to the next according to the instantaneous state of the world and the relative urgencies of the competing tasks, and the forced-scan implementation where attention is switched on a pre-programmed basis. The two hypotheses appeared to lead to nearly the same predicted performance effects when the duration of an individual glance away from the road was relatively short.

Visual and auditory tasks of comparable durations were compared for two hypothesized situations in terms of their effects on predicted steering performance. In both cases, the model predicted that the auditory implementation would be substantially less intrusive.

Guidelines for Further Model Application

There are basically three ways in which the model might be used: (1) calibration of model parameters via comparison of model and experimental output, (2) interpretation of experimental results, and (3) extension of the existing data base to situations not specifically explored. The latter category encompasses a host of potential applications, which includes experiment design; design and evaluation of a variety of systems, including those related to Intelligent Vehicle/Highway Systems (IVHS); and, of greatest relevance to the program sponsoring this work, the formulation of human factors guidelines for in-vehicle display design. The remainder of this discussion focuses on the extrapolative use of the model.

As with any mathematical model of the human operator, one has to be cautious in extending the model beyond known results because of the complexity of human response behavior. There is still much to be learned in terms of human information processing before the IDM can be fully calibrated. For example, the authors know of no theoretical models that relate quantitatively the performance on an auditory task to the fraction of attention allocated to that task. In general, one must rely on whatever human performance data exist to assign workload requirements to various auxiliary tasks, and use these as independent inputs to the model. For this and other reasons, the IDM is most reliably used to explore performance *trends*, rather than to predict absolute levels of driver/vehicle performance. That is, the model should be used to explore relative differences in performance across candidate systems, or from one task workload situation to another.

There are two primary modes of operation of the IDM, which are called here the controlled-scan and forced-scan implementation. The controlled-scan implementation requires the full capability of the IDM and is so called because the allocation of attention among competing tasks is assumed to be under the control of the driver and to be based on the driver's assessment of the situation at any given instant. The forced-scan implementation does not require the decision-making submodel and is used to explore the sensitivity of driving performance to specific attentional demand profiles assumed to be imposed by in-vehicle and other auxiliary tasks.

The IDM, suitably enhanced, has the potential to serve as the driver element in a variety of all-digital simulations in which micro-models of individual car/driver systems are desired. Candidate applications include the study of advanced traffic management systems (ATMS) and advanced vehicle control systems (AVCS). In the case of ATMS, the model could be used to explore the relative impact on driver behavior (and consequently, traffic flow) of voice and visual modes of presentation of on-route guidance information and commands. AVCS issues particularly amenable to model analysis include collision warning systems and transition between manual and automatic control, where system performance may be especially sensitive to attentional factors. The full IDM controlled-scan implementation would most likely be required in these cases because: (1) the timing of events would be largely determined as the simulation unfolds, rather than scripted ahead of time, and (2) driver decision-making capability would need to be modeled to ensure that the driver did not divert attention from the forward scene when it was not reasonable to do so.

The IDM is also required if the analyst specifically wishes to predict the attentional demands of candidate systems or of different driving environments.

For the IDM to be applied to simulations in which the driver reacts to the behavior of other vehicles, or to otherwise accommodate realistic driving situations, the model needs to be enhanced to include speed and headway control as well as the currently implemented constant-speed steering control. Such a modification would require a non-trivial implementational effort, but would not require a significant advance in the conceptual model.

The forced-scan implementation is recommended for purposes of sensitivity studies (including studies designed to formulate guidelines) where one is interested in the relative impacts of a particular set of attentional demand timelines, or in the effects of other task factors (e.g., road curvature) for one or more assumed task-loading conditions. For this type of application, the economy of model operation and the relative ease of data analysis will more than offset any slight decrease in simulation fidelity, especially where performance trends are of primary interest.

The forced-scan model may tend to yield substantially more pessimistic performance predictions than the controlled-scan implementation in situations where the driver attends to a highly demanding auditory task of long duration. Because the human driver retains the option to look at the forward scene while speaking or listening, cognitive attention to the forward scene can be increased whenever the driver detects an event that causes the driver to reassess the relative importance of attending to the driving and auditory tasks. Perhaps the best way to model this situation with the forced-scan model is to more closely mimic the typical visual-scan timeline in which the driver is allowed brief glances to the forward scene between scans to the auxiliary display. The auditory-task analog would be to assume intervals of limited time and high levels of attention to the auditory task separated by intervals of low levels of attention to the auditory task.

One issue that needs to be studied experimentally to allow more faithful modeling of the driver in a task-sharing environment is the effect of in-vehicle *control* activity. The discussion up to now has been devoted to auxiliary *monitoring* tasks that do not require the driver to remove a hand from the wheel. The assumption has been made that the effects of attentional diversion on steering performance are entirely due to the loss of information from the forward scene. When the driver performs an in-vehicle manual control function, however, the hand remaining on the steering wheel may not compensate perfectly for the imbalance of forces due to the removal of the hand operating the in-vehicle control, resulting in an unwanted control input that integrated over a substantial period of time, could have consequences overshadowing the effects of information loss. This factor can be readily accommodated with the IDM by a constant or stochastic wheel movement added to the intended wheel movement; experimental data are required, however, to quantify this process.

APPENDIX A. IMPLEMENTATION OF THE DRIVER/VEHICLE MODEL

OVERVIEW

Additional details regarding the methods and algorithms embodied in the driver/vehicle component of the integrated driver model (IDM) are presented. This material is intended for readers who have a working familiarity with linear, quadratic, Gaussian (LQG) optimal control and estimation theory.

Whereas the discussions of the driver/vehicle model formulation provided in chapters 1 and 2 were intended mainly to provide the reader with a conceptual framework, the material provided in this appendix is directed to the actual software implementation. Figure 23 shows the important elements and flow of information of the computerized driver model.

The model implementation as reflected in figure 23 differs in the following respects from what may have been implied in the previous discussions of the conceptual model:

- All dynamical response elements corresponding to the external task environment are implemented as a single dynamical system that includes not only the automobile response dynamics, but also any linear system elements needed to represent external disturbance or command inputs, sensor lags, etc.

- The optimal estimator is shown as being composed of three major elements: (1) a one-step predictor, (2) a set of internal display laws, and (3) estimator gains. The one-step predictor estimates the system state at time $t + \Delta$ on the basis of the estimated state at time t, where Δ is the simulation update interval. As discussed below, this prediction is accomplished using the driver's internal (mental) model of the system dynamics. The one-step prediction of system state is processed by the driver's internal model of the display laws to yield a one-step prediction of system outputs. The residual, which is the difference between the driver's noisy perceptual inputs and the one-step prediction of system outputs, is processed by the estimator gains to update the estimate of system state.

- Inclusion of the driver's control rate in the quadratic performance index dictates that the rate of change of control (e.g., steering wheel rotational velocity) be treated as the mathematical control input. The optimal control laws therefore yield a commanded control rate that is then integrated to yield the commanded control input corresponding to the driver's real-world control input (e.g., steering wheel displacement). To accommodate this treatment, the internal model of the system response dynamics used in the optimal estimator is augmented to include the commanded control as an additional system state as explained below.

- Because there will be many fewer control inputs (typically, one or two) than system outputs, time delay is implemented as being associated with the driver's control actions rather than with perceptual inputs in order to economize on storage and computational requirements. Because the model contains only linear dynamical response elements, the effects of delay on driver behavior and closed-loop system response are indifferent to the placement of delay within the control loop.

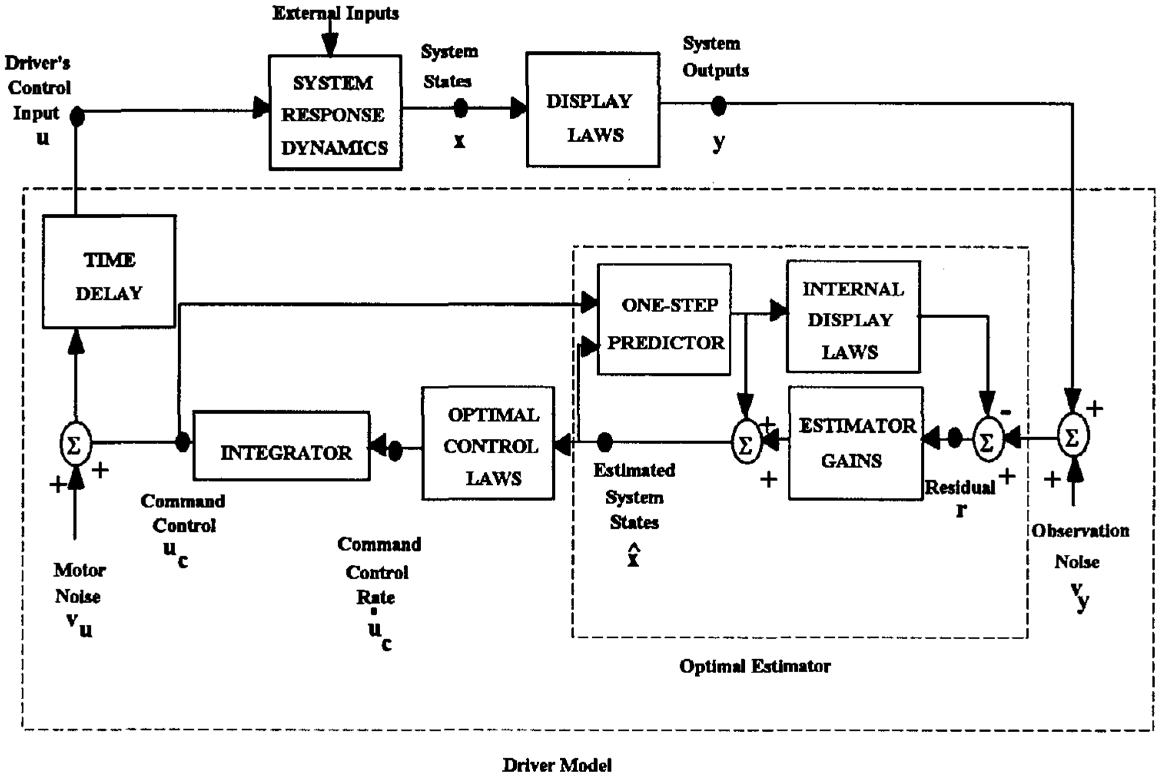

Figure 23. Implementation of the driver/vehicle model.

COMPUTATIONAL CONSIDERATIONS: THE EXTERNAL SYSTEM

The simulation model generates samples of predicted time histories for a fixed simulation interval specified by the user. The state variables describing the task environment (e.g., system response dynamics) are updated by the following transition matrix:

$$x(k+1) = e^{A\Delta}x(k) + \int_0^{\delta} e^{A\sigma}Bu(k)d\sigma \qquad (9)$$

where the Δ is the simulated update interval, **x** is the state vector for the external system, and **u** is the vector of inputs to the external system. In this context, the vector **u** includes all inputs to the external system, which, in general, will include external forcing functions as well as the driver's control inputs.

COMPUTATIONAL CONSIDERATIONS: THE DRIVER

We now consider some of the computational aspects of the portion of the model that relates to the driver's response strategy. In the following discussion, all references to system variables (e.g., state variables and matrices relating to response dynamics) refer to the driver's internal model of the system. In general, the internal model is equated to the model of the external system unless there are specific reasons for assuming otherwise.

Augmentation of the Problem Description

The vector of state variables associated with the driver's internalization of the control task is an augmented version of the state vector associated with the external world. First, because mathematical constraints require that the commanded control rate be considered as the control variable for the purposes of obtaining a problem solution, the state vector is augmented to include the true control input as a system state. Second, because the driver's internal model of system dynamics (unlike the external system model) uses a Pade approximation, further augmentation is needed to accommodate the additional state variable(s) required for this approximation.

A first-order Pade approximation is associated with each control variable. The general formulation for this dynamic system is:

$$y = \frac{s - 2/T}{s + 2/T} x \qquad (10)$$

where x is the undelayed variable, T is the amount of delay in seconds, y is the delayed variable, and s is the Laplace frequency variable.

The noise vector **w** must also be augmented to include the motor noise input $\mathbf{v_u}$. This system augmentation yields the following augmented system vectors and matrices associated with the driver's internal model of the control task, where **I** is an identity matrix having dimensionality equal to the number of driver control inputs:

$$\mathbf{x_o} = \begin{bmatrix} \mathbf{x} \\ \mathbf{u_d} \\ \mathbf{u} \end{bmatrix} \qquad \mathbf{w_o} = \begin{bmatrix} \mathbf{w} \\ \mathbf{v_u} \end{bmatrix}$$

$$A_o = \begin{bmatrix} A & B & -B \\ 0 & \frac{-2}{T}I & \frac{4}{T}I \\ 0 & 0 & 0 \end{bmatrix} \qquad \mathbf{B_o} = \begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \mathbf{I} \end{bmatrix}$$

$$\mathbf{E_o} = \begin{bmatrix} \mathbf{E} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \qquad \mathbf{C_o} = [\,\mathbf{C}\;\mathbf{D}\;\; \mathbf{-D}\,] \qquad (11)$$

The revised problem representation that is used for computing optimal estimation and control strategies is now

$$\dot{x}_O = A_O x_O + B_O u_C + E_O w_O$$

$$y_O = C_O x_O \qquad (12)$$

Additional details on system augmentation are given in Levison, Baron, and Junker (1976).

Optimal Control Laws

A two-step procedure is followed to compute the optimal control gains. First, control gains **L** appropriate to a continuous, statistically steady-state problem are computed as follows:

$$L = B_O K_O \tag{13}$$

where **K** is the solution of the Riccati equation

$$A_O' K_O + K_O A_O + Q_O - K_O B_O G^{-1} B_O' K_O \tag{14}$$

The matrices $\mathbf{Q_0}$ and **G** contain the coefficients related to the driver's cost function that determines the relative penalties of system errors and control activity.

A correction factor is then applied to convert these continuous gains to gains that are appropriate to the time-discrete formulation actually implemented as described by Kleinman, Baron, and Berliner (1977).

Note that the resulting time-discrete control gains are time-invariant gains appropriate for steady-state problems—they are not scheduled to truly optimize a time-varying control problem. The simulation model thus imbeds the assumption that the control (as opposed to estimation) aspect of the driving task can be treated as a piecewise-continuous control task, with potentially different constant gains being appropriate for different portions of the simulation. The gains can be varied in a step-wise fashion during the course of the problem solution by a redefinition of system matrices or of the cost weightings in the performance index.

Optimal Estimation

Because the driver model solves a one-step estimation problem, the state-estimation aspect of the driving task is solved in a time-varying fashion (even for simulations of steady-state tasks). The following computations relevant to driver response behavior are performed at each simulated update interval: (1) compute the Kalman estimator gains, (2) update the driver's estimate of the system state vector, (3) update the covariance of the driver's estimation error, and (4) compute the driver's control input.

To enhance readability we omit the subscript o in the equations presented in the remainder of this appendix. It is understood that all references to systems variables **x, A, B**, etc. refer to the augmented system as described previously.

The estimation problem solved in the driver model is that of minimizing the mean-squared error in the predicted state at time index k+1, given state and observational information at time index k. The discrete Kalman estimator gain is computed as

$$H = \Phi C \Sigma C' [C \Sigma C' + V]^{-1} \tag{15}$$

where **H** is the matrix of estimator gains, Φ is the transition matrix $e^{A\Delta}$, and Σ is the estimation error covariance matrix given information at time-index k. In order to provide a relatively stable observation noise covariance for computing the estimator gains, the observation noise samples (drawn from a random noise generator at each time index and scaled according to equation 4) are squared and then filtered with a first-order time constant of 0.5 s to compute the noise matrix **V**.

Once the filter gains have been computed, the state estimate is updated according to

$$x(k+1) = \Phi x + H(y + v_y - Cx) \tag{16}$$

where $\mathbf{x}(k+1)$ is the predicted state vector at time index k+1, given information at time index k. The quantities on the right side of the equation all reflect information at time index k. The variable x in this expression corresponds to the variable $\hat{x}$ of figure 23. The term $\Phi_{\mathbf{x}}$ is the one-step predictor mentioned previously, $\mathbf{H}$ is the matrix of estimator gains, and the expression ($\mathbf{y} + \mathbf{v_y} - \mathbf{Cx}$) is the residual (the difference between actual and expected perceptual inputs).

The computation shown in equation 16 applies when the driver is attending to the road. In the absence of roadway cues during visual attention to, say, an in-vehicle display, the $\mathbf{H}$ matrix is set to zero and the estimate is based solely on the driver's internal model of the situation.

Because the driver's internal model contains a Pade approximation to the time delay, which has been incorporated into the augmented state vector, no separate predictor beyond the one-step predictor described above is needed to deal with the delay.

Once the estimator gains have been computed, the estimation error covariance is updated as

$$\Sigma(k+1) = (\Phi - HC)\Sigma\Phi + EWE \tag{17}$$

Updating the Control Input

Because the driver model is a simulation model operating at discrete update intervals, the question arises as to what value of driver control input should be used as an input to the external system—the value appropriate to the start of the update interval or the value appropriate to the end of the update interval. We compromise by using the average of the two values.

We first compute the commanded control rate as

$$\dot{u}_c = -Lx \tag{18}$$

where $\mathbf{L}$ is the optimal control gain matrix, corrected for the discrete-time formulation. The minus sign is used to conform with standard control-theory practice for designating negative (i.e., restorative) feedback in a closed-loop control system.

Let $\mathbf{u_{c1}}$ be the commanded control computed for the end of time index k. The commanded control computed for the end of time index k+1 is a simple trapezoidal integration of the previous commanded control input:

$$u_{c1}(k+1) = u_{c1} + \Delta\dot{u}_c \tag{19}$$

To compute the pilot's control input $\mathbf{u_c}$ for the next sample, we average the initial and final commanded control and add a sample of motor noise, if relevant:

$$u_c = 0.5(u_{c1} + u_{c1}(k+1)) + v_u \tag{20}$$

Although delay is represented by a Pade equivalent in the driver's internal model, a more veridical treatment of delay is necessary to maintain proper time relationships of the predicted time histories of the external system state variables. Therefore, a buffer of a length consistent with the assumed delay is inserted between the control commanded by the driver and the external system. This buffering operation is represented by the time delay element shown in figure 23.

Appendix B. Specification of Non-Driving Tasks

OVERVIEW

This appendix reviews the portion of the integrated driver model (IDM) procedural component concerning the representation of non-driving tasks, with emphasis on the format used in specifying new tasks. Recall that this model component is currently written in Lisp; hence, references to and examples of Lisp structures appear in this discussion. Following this overview is a description of the standardized format for specifying non-driving tasks, followed by three illustrative examples.

Figure 24 provides a diagram of the information flow concerning the enabling of tasks. The block labeled "All Tasks" represents the entire pool of tasks that are resident in the IDM at any given configuration. All tasks other than the driving (steering) tasks are specified as described in this section. As the simulation proceeds, individual tasks are then selected via the scripting mechanism and are added to a set of Currently Active Tasks (the set of tasks that are currently capable of being executed if the conditions for their execution are met).

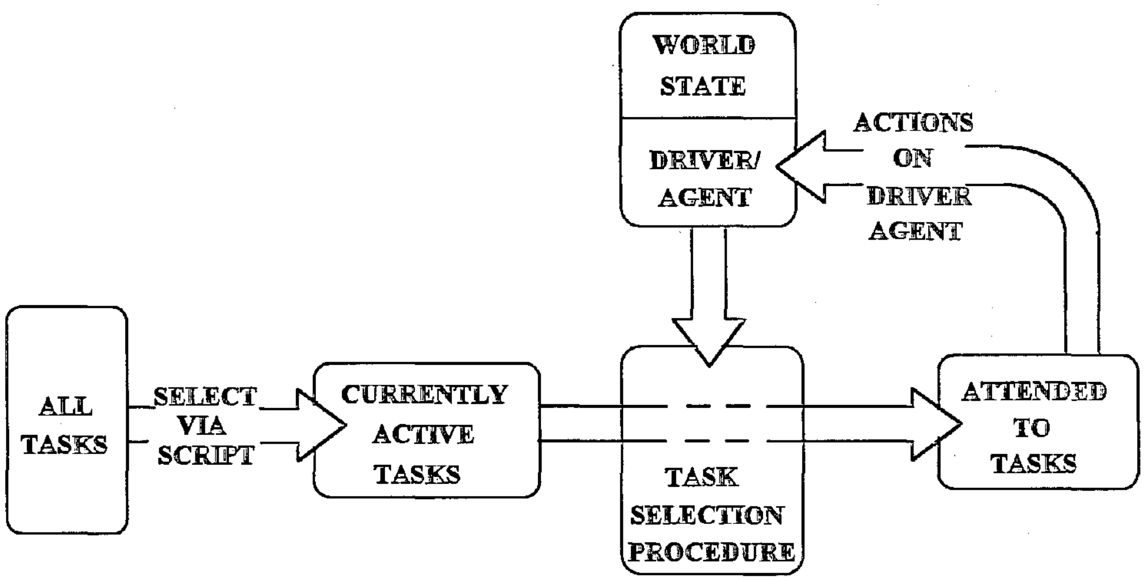

Figure 24. Flow of task information.

On each tick, each task in this set of Currently Active Tasks is examined by the underlying Lisp program. Combined with this is the information concerning the current state of the driver-agent and of the external world. (This knowledge of the state of the external world is obtained through the driver-agent and his model and memory of the external world.) At this point, the Lisp program functions as a filter to select those tasks whose conditions are such that they can be executed on this task. If multiple tasks can be run at a single time, the rules outlined elsewhere in this report are invoked to uniquely determine the task(s) that are to be executed.

The selected set of tasks are then executed and they, in turn, affect the state of the driver-agent in terms of perceptual information and requirements for response by the driver. (Note that these actions affecting the driver-agent can, in turn, affect the state of the Lisp world model, e.g., through button-presses, changes to parameters, etc.) These actions are for the primary low-level interface between the tasks and the underlying Lisp model of the driver-agent.

Specification of Tasks

Two levels of programming are required to model a new non-driving task. What we might consider as the high-level model is a standardized Lisp form that can be read into the model much as one would read in a numerical parameter. This standardized format, which is described below, may be considered as a specialized language for describing tasks.

Low-level coding is also required to interface the actions represented in the outer layer to the world and agent models (i.e., to pass information to and from the rest of the IDM). This code must represent a specific interaction with the underlying Lisp model and therefore cannot be constrained to the same rigid structure as presented below for high-level coding.

There is typically a many-to-one relationship between the high and low levels of code in that a relatively small number of low-level functions can be expected to support a much larger number of high-level models. As we show below, the low-level function related to pressing a button supports more than one higher level activity.

The high-level specification of a task has the following general form:

```
(defccmtask <TASK-NAME> <TASK-ARGUMENTS>
     <TASK-TYPE>
 <TASK-PARAMETER-KEY-1> <TASK-PARAMETER-VALUE-1>
 <TASK-PARAMETER-KEY-2> <TASK-PARAMETER-VALUE-2>
 ....
```

<TASK-NAME>
The name of the task. This name should be unique across the set of all available tasks.

<TASK-ARGUMENTS>
A list of arguments that is local to the task.

<TASK-TYPE>
A list of one or more elements specifying the type of task. Currently implemented task-types are:

> TOP-LEVEL-TASK:
> This is the primary type of main task to be accomplished by the operator. An example of this type of task would be the activity of turning on and dialing the phone.

A given top-level task typically divides into a number of sub-tasks. These types of tasks, together with the TOP-LEVEL-TASK, specify a number of "child" tasks to be run, i.e., they form the nodes in the tree-like task hierarchy. These are of the following types:

> SEQUENTIAL-TASK:
> This type of task performs a number of sub-tasks in sequential order and then terminates.

ROTATION-TASK:
This type of task performs a sequence of sub-tasks in order, and then repeats the sequence. The sequence is repeated until some specified termination condition is met.

The lowest level tasks have the following form (these tasks essentially form the "leafs" in the tree-like task hierarchy):

FIXED-DURATION-TASK:
This type of task performs for a fixed, given duration and then terminates.

SIMPLE-TASK:
This type of task continues to operate until a specified termination condition is met.

<TASK-PARAMETER-KEY-n> <TASK-PARAMETER-VALUE-n>

The task parameters specify certain characteristics and behavior for a given task. They are specified by a sequence of keyword/value pairs. The following parameters can be specified:

FIXED-DURATION-FORM:
Used by a FIXED-DURATION-TASK, this form is evaluated to specify the duration during which the task operates.

TASK-FORMS:
This is a list specifying the child tasks that a TOP-LEVEL-TASK, a SEQUENTIAL-TASK, or a ROTATION-TASK can specify.

The specification for each task in the TASK FORMS list can be in one of two forms:

1. <TASK-NAME>
   In this instance, the task takes no arguments.

2. (<TASK-NAME> <TASK-ARG-KEYNAME-1> <TASK-ARG-VALUE-1> ...)
   In this case, the specification for the task consists of a list containing the name of the task followed by a sequence of keyname/value of the initialization arguments for the task.

TICK-FORM:
This form is evaluated once for each simulation time-step for which the task is active. This feature typically functions as the primary low-level interface between the tasks and the underlying Lisp model of the driver and his state.

LOCKUP-DURATION-FORM:
This form is evaluated and specifies the duration during which the task may not be interrupted following its initial creation.

VISUAL-RESOURCE-FORM:
This is a symbol that specifies the focal-location of the eyes as required by this task.

PMOTOR-RESOURCE-FORM:
This is a symbol that specifies the location of the (right) hand as required by this task.

SUPPRESS-PMOTOR-RESOURCE-WAIT?:

If this value is TRUE, then this task does not wait for the PMOTOR resource (i.e., the right hand) to reach its required location before the task can be performed. This is used, for example, by the top-level driving task to specify that driving can take place without the right hand (i.e., because the left hand is always on the steering wheel). (The default value is NIL, i.e., the wait is not suppressed.)

NO-RUN-LOSS-FUNCTION:
For top-level tasks, this form specifies the loss that would be accrued if this task were not to be run for a given tick.

PRIMARY-TASK-TYPE:
The primary psychomotor resource that is used by this task. Currently specified values are:
visual
auditory

SECONDARY-TASK-TYPE:
The secondary cognitive resource that is used by this task. Currently specified values are:
spatial
cognitive

NON-INTERRUPTIBLE?:
If this value is TRUE, then this task (and its children, if any) cannot be interrupted once it has begun executing. (This is typically used only for low-level, compact tasks to specify behavior that is procedurally assumed to be an indivisible unit.) The default value is NIL, i.e., the task is typically interruptible.

TERMINATED?-FORM:
This form is evaluated on each tick. If it returns TRUE, then the execution of this task is terminated.

TERMINATE-ON-INTERRUPT?:
If this value is TRUE, then the execution of this task will be terminated if the task is interrupted.

TAIL-TASK-FORMS:
This form is similar to the TASK-FORMS specified above, except that the task specified by this form is spawned upon the completion of the calling task. This is used to sequentially launch independent top-level tasks.

Note: Some of the above forms are evaluated by the task at run time. They are evaluated in an environment in which the values of a number of useful global variables are set. Primary among these are the variables:
AGENT : The agent-object (the driver) that is running the current task.
TASK : The current task itself.

EXAMPLES OF SPECIFIC TASKS

Below are examples of three specific tasks: (1) the task PRESS-BUTTON, a low-level task responsible for directing the driver-agent to press a specified button; (2) READ-HELPER-INTERNAL, a mid-level task that is used when the agent reads a section of text; and (3) END-

PHONE-CONVERSATION-TOP-LEVEL, a top-level task that is responsible for terminating a telephone conversation.

Task: PRESS-BUTTON

This task is expressed as the following Lisp form:

```
(defccmtask PRESS-BUTTON (BUTTON-ID)
         (FIXED-DURATION-TASK)
     FIXED-DURATION-FORM
          0.2
     :TICK-FORM
        '(PRESS-BUTTON-FUNCTION *AGENT* (slot-value *TASK*
        'BUTTON-ID))
     :VISUAL-RESOURCE-FORM
          '(slot-value *TASK* 'BUTTON-ID)
     :PMOTOR-RESOURCE-FORM
          '(slot-value *TASK* 'BUTTON-ID)
     :NON-INTERRUPTIBLE?
          t
```

PRESS-BUTTON is the low-level task that is called when the driver-agent presses a button or key. It is of the type FIXED-DURATION-TASK. This task has the argument:

BUTTON-ID
This is a parameter that is passed into the task when it is called that specifies the identification of the button or key to be pressed.

The following parameters are also specified for this task:

FIXED-DURATION-FORM:
This is the amount of time for which the task will run.

TICK-FORM:
For each tick that the task is active, the function PRESS-BUTTON-FUNCTION is run. It is responsible for directing the driver to press the indicated button in the (Lisp) model of the car. (As mentioned above, the TICK-FORM typically functions as the primary low-level interface between the tasks and the underlying Lisp model of the driver and his state.)

The function PRESS-BUTTON-FUNCTION is a function with two arguments. The first argument is *AGENT*, i.e., the agent-object. The second argument is the ID of the button to be pushed. This value is obtained by examining the slot value named BUTTON-ID in the calling task object. (As discussed at the beginning of appendix B, *AGENT* and *TASK* are special variables that are bound to the appropriate value in the environment in which the tick form is evaluated.)

VISUAL-RESOURCE-FORM:
This specifies location of the eyes necessary for this task to execute. The value is obtained by examining the value stored in the task's local argument BUTTON-ID.

PMOTOR-RESOURCE-FORM:
This specifies location of the right hand necessary for this task to execute. The value is obtained by examining the value stored in the task's local argument BUTTON-ID.

NON-INTERRUPTIBLE?:
The task is specified to be uninterruptible for the duration of its execution.

Task: READ-HELPER-INTERNAL

The expression for this task is

```
(defccmtask READ-HELPER-INTERNAL (READ-LOCATION
                     (CURRENT-NUMBER-WORDS-READ :initform 0))
          (SIMPLE-TASK)
     :TICK-FORM
     '(READ-HELPER-INTERNAL-FUNCTION *AGENT* *TASK*)
     :LOCKUP-DURATION-FORM
          '(human-parameter *AGENT* 'READ-MONITOR-TCOMMIT)
     :TERMINATED?-FORM
          '(<= (word-contents *AGENT*
                         (slot-value *TASK* 'READ-LOCATION))
          (apply-parent-task *TASK* 'TOTAL-NUMBER-WORDS-READ))
     :VISUAL-RESOURCE-FORM
          '(slot-value *TASK* 'READ-LOCATION)
```

READ-HELPER-INTERNAL is the main, low-level task that is executed when the driver-agent reads text (from, for example, the display screen). It is of type SIMPLE-TASK, and will continue executing until it achieves a specific TERMINATION condition.

This task has two arguments:

READ-LOCATION:
This is a parameter that is passed into the task when it is called in order to specify the location of the text that is to be read.

CURRENT-NUMBER-WORDS-READ:
This functions as an internal variable that keeps track of the number of words that have been read up until the present time.

The following parameters are also specified for this task:

TICK-FORM:
For each tick that the task is active, the function READ-HELPER-INTERNAL-FUNCTION is run. It is responsible for updating the driver's state to represent the reading that occurs during one tick of the simulation. Among other things, it is responsible for updating the value of the local argument CURRENT-NUMBER-WORDS-READ described above. (As mentioned above, the :TICK-FORM typically functions as the primary low-level interface between the tasks and the underlying Lisp model of the driver and his state.)

LOCKUP-DURATION-FORM:

This specifies the amount of time after which the task is created during which it must run without being interrupted. The value is obtained by querying the driver-agent for the value of the run-time parameter READ-MONITOR-TCOMMIT that has previously been set by the analyst.

TERMINATED?-FORM:
This form is evaluated on each tick to determine if the conditions for terminating the task have been met. For this task, these conditions will be met when the TOTAL-NUMBER-WORDS-READ in the current task context is equal to the number of words currently at the READ-LOCATION.

VISUAL-RESOURCE-FORM:

This specifies location of the eyes necessary for this task to execute. The value is obtained by examining the value stored in the task's local argument READ-LOCATION.

Task: END-PHONE-CONVERSATION-TOP-LEVEL

The top-level code for the final example is:

```
(defccmtask END-PHONE-CONVERSATION-TOP-LEVEL ()
     (TOP-LEVEL-TASK)
     :TASK-FORMS
          '((PRESS-BUTTON :BUTTON-ID :PHONE-END)
          SETDOWN-HANDSET)
     :NO-RUN-LOSS-FUNCTION
          'END-PHONE-CONVERSATION-TOP-LEVEL-LOSS-FUNCTION
     :PRIMARY-TASK-TYPE
          :visual
```

END-PHONE-CONVERSATION-TOP-LEVEL is an example of a top-level task. It is called by the driver-agent to end a phone conversation. The following parameters are specified for this task:

TASK-FORMS:
This top-level task will sequentially spawn two subtasks:

1. The Task PRESS-BUTTON with the argument value :BUTTON-ID being set to an initial value of :PHONE-END. (This task is responsible for directing the drive-agent to press the PHONE-END button.)
2. SETDOWN-HANDSET, a task of no arguments that is responsible for hanging up the handset.

NO-RUN-LOSS-FUNCTION:
The function END-PHONE-CONVERSATION-TOP-LEVEL-LOSS-FUNCTION is run to determine the penalty associated with not executing this task (see above).

PRIMARY-TASK-TYPE:
This top-level task is specified to be of the VISUAL type.

References

Baron, S. and Levison, W.H. (1975). "An Optimal Control Methodology for Analyzing the Effects of Display Parameters on Performance and Workload in Manual Flight Control." *IEEE Trans. on Systems, Man, and Cybernetics*, 5(4), pp. 423-430.

Baron, S. and Levison, W.H. (1977). "Display Analysis With the Optimal Control Model of the Human Operator," *Human Factors*, 19(5), pp. 437-457.

Becker, W. and Fuchs, A.F. (1969). "Further Properties of the Human Saccadic System: Eye Movement and Correction Saccades With and Without Visual Fixation Points," *Vision Research*, 9, pp. 1247-1258.

Card, S.K., English, W.K., and Burr, B.J. (1978). "Evaluation of Mouse, Rate-Controlled Isometric Joystick, Step Keys, and Test Keys for Text Selection on a CRT," *Ergonomics 21*, pp. 601-608.

Card, S.K., Moran, T.P., and Newell, A. (1980). "The Keystroke-Level Model for User Performance Time With Interactive Systems," *Comm ACM*, 2, pp. 396-408.

Corker, K., Cramer, N.L., and Henry, E.H. (1990). *Methodology for Evaluation of Automation Impacts on Command Control C2 Systems: Implementation* (BBN Report No. 7242). Wright-Patterson AFB, OH: Armstrong Laboratory, Logistics Research Division.

Gibson, E.J. and Levin, H. (1975). *The Psychology of Reading*. Cambridge, MA: MIT Press.

Glanzer, M., Dorfman, D., and Kaplan, B. (1981). "Short-Term Storage in the Processing of Text," *Journal of Verbal Learning and Verbal Behavior*, 20, pp. 656-670.

Godthelp, H., Milgram, P., and Blaauw, G.J. (1984). "The Development of a Time-Related Measure to Describe Driving Strategy," *Human Factors*, 26, pp. 257-268.

Green, P. (1993). *Measures and Methods Used to Assess the Safety and Usability of Driver Information Systems* (Technical Report No. UMTRI-93-12). Ann Arbor, MI: University of Michigan Transportation Research Institute.

Green, P., Hoekstra, E., Williams, M., Wen, C., and George, K. (1993). *On-the-Road Testing of an Integrated Driver Information System* (Technical Report No. UMTRI-93-32). Ann Arbor, MI: University of Michigan Transportation Research Institute.

Green, P. and Olson, A. (1989). *The Development and Use of the UMTRI Driving Simulator* (Technical Report No. UMTRI-89-25). Ann Arbor, MI: University of Michigan Transportation Research Institute.

Green, P., Williams, M., Serafin, C., and Paelke, G. (1991). "Human Factors Research on Future Automotive Instrumentation: A Progress Report." In *Proceedings of the Human Factors Society 35th Annual Meeting*, pp. 1120-1124. Santa Monica, CA: Human Factors Society.

Hess, R.A. and Wheat, I.W. (1976). "A Model-Based Analysis of a Display for Helicopter Landing Approach," *IEEE Trans. on Systems, Man and Cybernetics*, 6(7), pp. 505-511.

Hoekstra, E., Wen, C., Williams, M., George, K., and Green, P. (1993). *Effects of Information Modality and Landmarks on Route Guidance/Traffic Information System Usability* (Technical Report No. UMTRI-93-31). Ann Arbor, MI: University of Michigan Transportation Research Institute.

Kleinman, D.L., Baron, S., and Levison, W.H. (1970). "An Optimal-Control Model of Human Response, Part I: Theory and Validation," *Automatica*, 6, pp. 357-369.

Kleinman, D.L., Baron, S., and Levison, W.H. (1971). "A Control Theoretic Approach to Manned-Vehicle Systems Analysis," *IEEE Trans. on Auto. Control*, 16(6), pp. 824-833.

Levison, W.H. (1971). *The Effects of Display Gain and Signal Bandwidth on Human Controller Remnant* (Technical Report No. AMRL-TR-70-93). Wright-Patterson Air Force Base, Ohio: Aerospace Medical Research Laboratory.

Levison, W.H. (1978). *Analysis and In-Simulator Evaluation of Display and Control Concepts for a Terminal Configured Vehicle in Final Approach in a Windshear Environment* (Technical Report No. NASA CR-3034). Hampton, VA: NASA Langley Research Center.

Levison, W.H. (1979). "A Model for Mental Workload in Tasks Requiring Continuous Information Processing." In N. Moray (Ed.), *Mental Workload: Its Theory and Measurement.* New York and London: Plenum Press.

Levison, W.H. (1983). *Development of a Model for Human Operation Learning in Continuous Estimation and Control Tasks* (Report No. AFAMRL-TR-83-088). Wright-Patterson Air Force Base, Ohio: Air Force Aerospace Medical Research Laboratory.

Levison, W.H. (1985). *Application of the Optimal Control Model to the Design of Flight Simulation Experiments* (Society of Automotive Engineers Paper No. 851903). Presented at the Aerospace Technology Conference and Exposition, CA.

Levison, W.H. (1989). "Applications of Human Performance Models to System Design." In G.R. McMillan, et al. (Eds.), *The Optimal Control Model for Manually Controlled Systems,* pp. 185-198, New York: Plenum Press.

Levison, W.H., Baron, S., and Kleinman, D.L. (1969). "A Model for Human Controller Remnant," *IEEE Trans. on Man-Machine Systems,* 10(4).

Levison, W.H., Elkind, J.I., and Ward, J.L. (1971). *Studies of Multivariable Manual Control Systems: A Model for Task Interference* (Technical Report No. NASA CR-1746). Moffett Field, CA: NASA Ames Research Center.

Levison, W.H. and Rickard, W. (1981). "Closed-Loop Criteria for Assessing Longitudinal-Axis Handling Qualities of Transports in Final Approach." In *Proceedings of AIAA Guidance and Control Conference,* pp. 470-479. Albuquerque, NM: American Institute of Aeronautics and Astronautics.

Munger, S.J., Smith, R.W., and Payne, D. (1962). *An Index of Electronic Equipment Operability: Data Store* (Report No. AIR-C43-1/62-RP[1]). Pittsburgh, PA: American Institute for Research.

Senders, J.W., Kristofferson, A.B., Levison, W.H., Dietrich, C.W., and Ward, J.L. (1967). *The Attentional Demand of Automobile Driving* (BBN Technical Report No. 1482), Cambridge, MA: Bolt Beranek and Newman, Inc.

Serafin, C., Wen, C., Paelke, G., and Green, P. (1993). *Development and Human Factors Tests of Car Phones* (Technical Report No. UMTRI-93-17). Ann Arbor, MI: University of Michigan Transportation Research Institute.

Simpson, C.A. and Marchionda-Frost, K. (1983). "Synthesized Speech Rate and Pitch Effects on Intelligibility of Warning Messages for Pilots." *Human Factors*, 26, pp. 509-518.

Teichner, W.H. and Krebs, M.J. (1972). "Laws of the Simple Visual Reaction Time," *Psychological Review*, 79, pp. 344-358.

Wickens, C.D. and Liu, Y. (1988). "Codes and Modalities in Multiple Resources: A Success and a Qualification," *Human Factors*, 30, pp. 599-616.

Zwahlen, H.T. and Balasubramanian, K.N. (1974). "A Theoretical and Experimental Investigation of Automobile Path Deviations When Driver Steers with No Visual Input," *Transportation Research Record #520*, pp. 25-37.

Zwahlen, H.T. and DeBald, D.P. (1986). "Safety Aspects of Sophisticated In-Vehicle Information Displays and Controls." In *Proceedings of the Human Factors Society 30th Annual Meeting, 1986,* pp. 256-260. Santa Monica, CA: Human Factors Society.