Brake Initiation and Braking Dynamics: A Human-Centered Study of Desired ACC Characteristics

Michael A. Goodrich, Erwin R. Boer and Hideaki Inoue

CBR TR 98-5

Abstract

The driver interprets and responds to sensory input according to the context provided by a mental model—an internal representation employed to encode, predict, and evaluate the consequences of perceived and intended changes to the operator's current state within the dynamic environment. To emulate driver behavior, we develop a multiple mental model framework that uses rule-based task switching to coordinate multiple skill-based controllers. We employ satisficing decision theory (SDT) to emulate rule-based task switching, and model predictive control (MPC) to emulate skill-based performance execution.

Decision makers in naturalistic settings employ moderation in generating behavior. SDT, which compares a benefit-like attribute called *accuracy* against a cost-like attribute called *rejectability*, is a mathematical realization of the notion of moderation. Accuracy and rejectability are represented as fuzzy set membership functions; such set membership functions, identified from experiments in longitudinal control, efficiently partition perceptual space into conditions which generate active braking and conditions which permit nominal behavior. In words, these set membership functions embody the intuition of "expedient but safe" moderate behavior.

Given the decision to brake, braking dynamics are characterized by smooth trajectories in perceptual space terminating at infinite time to collision and desired time headway. MPC, which determines control by evaluating consequences over a receding planning horizon, is a method that can emulate these braking dynamics. MPC parameterizes this trajectory in terms of a weighted perceptual distance from a target state balanced by the cost of control. This parameterization generates trajectories that closely match observations, but exhibits some sensitivity to initial perceptual conditions.

mike@cs.byu.edu

Cambridge Basic Research
4 Cambridge Center, Cambridge, MA 02142 USA
Tel: (617) 374-9650 Fax: (617) 374-9697

December 1, 1998

1 Introduction and Notation

As our part of an ongoing effort to understand driver interaction with advanced vehicle safety systems, we have examined the relationship between observed driver behavior and acceptable automated longitudinal systems. Our objective has been to construct mathematical models of automobile driver longitudinal control that not only emulate driver behavior, but also enhance our understanding of the processes that generate such behavior. A mathematical model that can emulate longitudinal driver control must avoid collisions, regulate vehicle speed about a desired value, maintain safe time headway, and produce comfortable acceleration and deceleration dynamics. The design of such a model is constrained by perceptually available cues, by existing modeling methods, and by human factors considerations. We construct a model of human car-following behavior that uses satisficing decision theory (SDT) to switch between skill-based tasks, and model predictive control (MPC) to execute these skill-based control tasks. Such a driver model not only identifies the perceptual cues that trigger active braking, but also emulates the resulting perception-driven deceleration dynamics. This model may contribute to intelligent vehicle system design (specifically ACC design), especially when human and automation share responsibility. In presenting this model, our intention is not to exhaustively specify and test a model of human cognition. Instead, we desire to create a computational model of behavior with not only a high degree of prediction accuracy but also a substantial amount of explanatory power. Our approach is to create a model motivated by justifiable experimental support and blended with reasonable modeling assumptions.

1.1 Multiple Mental Model Framework

Many aspects of cognitive decision-making have been described in terms of mental models [21, 36]. A mental model is an internal representation employed to encode, predict, and evaluate the consequences of perceived and intended changes to the operator's current state within the dynamic environment. We define a mental model $\mathcal{M}$ as a triplet consisting of the perceived state of the environment Θ, a set of decisions or actions U, and a set of ordered consequences C that result from choosing $u \in U$ when $\theta \in \Theta$ obtains. According to this specification, a mental model not only encodes the relation between the input-action pair[1] (θ, u) and the predicted consequence c, but also induces an evaluation of preferences among consequences (see Figure 1, and compare to related figures in [2, 34, 46]). In other words, the mental model $\mathcal{M}$ provides the context for meaningfully interpreting sensory information and generating purposeful behavior.

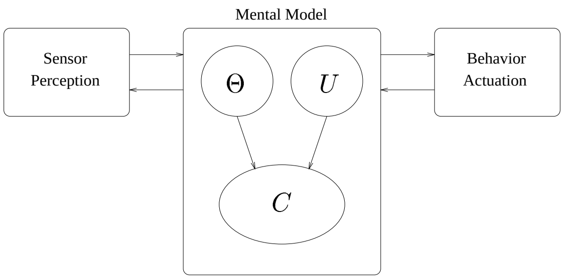

Figure 1. Working specification of a mental model.

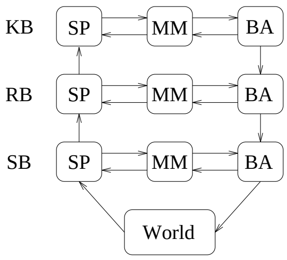

Figure 2. Interaction within a society of mental model agents. SP=sensor perception, MM=mental model, and BA=behavior actuation.

In driving, human cognition can be described using multiple mental models (treated as agents) that can be organized into a society of interacting agents [16, 36, 42]. Coordination within this society not only determines which agents contribute to driver behavior, but also which agents can employ attentional resources. A multi-resolutional society of interacting mental models organized into a three level hierarchical structure (see Figures 2 and 4) can be constructed corresponding to Rasmussen's knowledge-based (KB), rule-based (RB), and skill-based (SB) behaviors[2] [8, 38, 40, 41, 46]. At the KB level of this hierarchy, the agent role is supervisory; at the RB level, the agent role is task management; and at the SB level, the agent role is task execution. Intuitively, the KB, RB, and SB agents think, monitor, and control, respectively[3].

[1] A decision u is often treated as a mapping from Θ into the set of consequences[13].

[2] These layers also appear to correspond to Saridis *organization*, *coordination*, and *execution* levels, respectively, for intelligent machine design [44].

[3] Ramifications of this structure is that much of SB behavior is not consciously processed but instead proceeds in a closed

Each mental model $\mathcal{M}$ will be described as being enabled/disabled and engaged/disengaged (see Table 1. When $\mathcal{M}$ is *enabled* the mental model is actively influencing human behavior generation, and when *disabled* the mental model has no direct influence upon behavior. When *engaged* the mental model holds attention whereby environmental information is actively perceived and interpreted, and when *disengaged* the mental model releases attention whence no such active perception occurs. In terms of Figures 1-2, the mental model

		BEHAVIOR GENERATION	
		Enabled	Disabled
SENSORY PERCEPTION	Engaged	interpret and act	observe without acting
	Disengaged	act without observing	off

Table 1. Activity states of mental models.

is enabled if the arcs between the mental model and behavior/actuation are active (whence behavior u is actuated) and the mental model is engaged if the arcs between the mental model and sensor/perception are active (whence θ is actively perceived). We suppose that $\mathcal{M}$ need not be enabled to be engaged, nor conversely. We develop a structure to manage which mental models contribute to behavior generation, and present some preliminary ideas about which mental models consume attentional resources. These mental model agents operate within the context of overall complex human behavior.

1.2 Problem Description and Notation

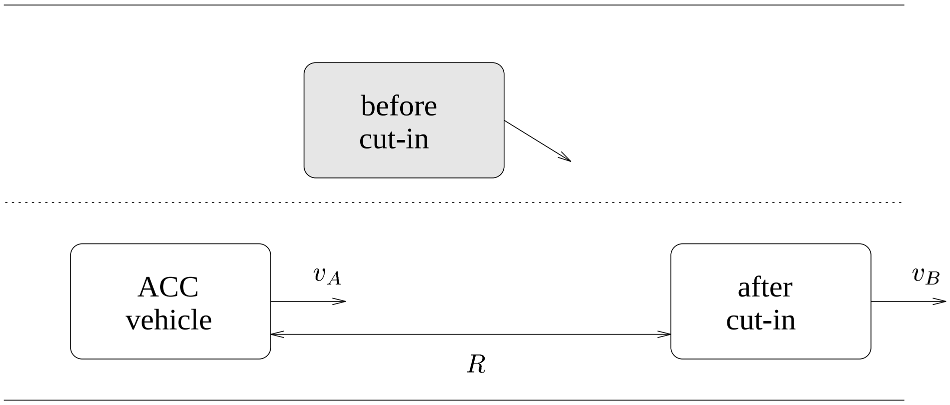

Figure 3. "Cutting in" problem. The cut-in vehicle prior and subsequent to the cut-in event is represented by a shaded box and an open box, respectively.

To determine SDT-based models of driver behavior, we will focus on the "cutting in" problem wherein vehicle B cuts in front of the driver's vehicle (vehicle A) as diagrammed in Figure 3. Subsequent to a cut-in event, we refer to the cut-in vehicle as the lead vehicle. In the figure, v_A and v_B represent the velocities of the driver's vehicle and the cut-in (lead) vehicle, respectively, $v_R = v_A - v_B$ represents the relative velocity[4] between the vehicles, and R represents the range (relative distance) between the vehicles. From these variables, we construct a state vector $\boldsymbol{x} = [v_A, R, v_R]^T$ that, depending upon the accelerations of vehicles A and B (denoted u_A and u_B, respectively) and assuming disturbance free dynamics, yields a discrete time dynamical system to describe how the state $\boldsymbol{x}$ changes over time (indexed by k) $\boldsymbol{x}_{k+1} = \boldsymbol{g}(\boldsymbol{x}_k, u_A, u_B)$. By convention, $u_A < 0$ indicates that the driver is pressing the brake pedal, and $u_A > 0$ indicates that the driver is pressing the accelerator pedal. The same convention applies for u_B.

Shifting focus from a world centered perspective to a driver centered perspective, we construct a model of car following behavior using a discrete time dynamical state space representation that possesses the following five desirable features:

perception-response loop [8].

[4]Note that in these definitions, we violate the usual sign convention and define $v_R = -\frac{dR}{dt}$. We do this so for convenience in defining time to collision.

Feature 1: state variables $\boldsymbol{\chi}$, possibly different from $\boldsymbol{x}$, are perceivable by driver,

Feature 2: the space spanned by $\boldsymbol{\chi}$ (denoted $\text{sp}(\boldsymbol{\chi})$) equals $\text{sp}(\boldsymbol{x})$,

Feature 3: an internal dynamical model of perceptual state transitions $\boldsymbol{\chi}(k+1) = \boldsymbol{f}(\boldsymbol{\chi}(k), u_A, u_B)$ can be constructed ($\boldsymbol{f}$ denotes the dynamical response in space $\boldsymbol{\chi}$, and $\boldsymbol{g}$ denotes the related dynamical response in space $\boldsymbol{x}$),

Feature 4: a control law $u_A = \pi(\boldsymbol{f}, \boldsymbol{\chi})$ can be constructed from the internal model and the observed perceptual state using cognitively plausible decision mechanisms, and

Feature 5: decision planes can be described in a low dimensional subspace of $\text{sp}(\boldsymbol{\chi})$ (i.e., decisions depend on relatively few variables).

These five desirable features are motivated by the multiple mental model framework. Clearly, SB task execution requires perception of cognitively feasible and ecologically informative cues (Features 1-2), and can employ an internal model structure to effect behavior (Features 3-4). Coordination of SB behaviors is performed by RB task managers who employ perceptual decision planes to manage task switching (Feature 5).

In constructing $\boldsymbol{\chi}$, we consider time headway and inverse time to collision, respectively defined as follows[5]:

$$T_h = \frac{R}{v_A} \quad T_c^{-1} = \frac{v_R}{R}, \tag{1}$$

because there exists evidence that these perceptual cues can be directly perceived (Feature 1) by people (see, for example, [24, 56]). Given these perceptual values, a perceptual state can be defined as $\boldsymbol{\chi} = [T_c^{-1}, T_h, v_A]^T$. Note the one-to-one (except on the surface $v_A = 0$) and onto mapping from the physical state space $\boldsymbol{x} = [R, v_A, v_R]^T$ to the perceptual state space $\boldsymbol{\chi}$ whence $\text{sp}(\boldsymbol{\chi}) \approx \text{sp}(\boldsymbol{x})$ (Feature 2). From an internal model (Feature 3), the driver can form estimates of future perceptual states ($\ell \Delta t$ seconds into the future) yielding predictions $\hat{\boldsymbol{\chi}}(k+\ell)$ which can be used to generate behavior (Feature 4) and manage the agent society (Feature 5).

1.3 Literature Review

Model predictive control (MPC), also known as moving horizon and receding horizon control, is a method for designing controllers that operate in nonlinear, constrained, and uncertain environments. MPC design requires the identification of a system model, and the specification of a system performance metric defined over a finite planning horizon. The success of MPC is due, in large part, to its ability to handle uncertain nonlinear systems with state and input constraints such as those found in complex industrial processes. Successes in application are supported by theoretical advances, such as the characterization and specification of sufficient conditions for stability and observability [30, 35, 43, 45, 50].

We adopt the satisficing approach described in [15, 18]. This approach represents the consequences of a decision by a cost-like attribute called *rejectability* and a benefit-like attribute called *accuracy.* Partitioning the consequence set into these attributes recalls the generalized potential field (GPF) approach to path planning and obstacle avoidance [19, 37]. In the GPF methodology, a goal is represented as an attractive potential, obstacles are represented as repulsive potentials, and the path along the negative gradient of the combined potentials is selected as a collision free path. With the application of harmonic potential fields [9, 22], the problem of a robot remaining in an attractive local minima is avoided, but the problems with forming a globally attractive potential field in the presence of moving obstacles remains. Additionally, although computationally efficient, GPF's do not consider the optimality of the resulting path [52]. A method proposed in [37] deals with the moving obstacle problem using a GPF formalism by incorporating the *view-time* concept. This concept appears to be a special case of a receding planning horizon as employed in MPC. By employing fuzzy representations of cost and benefit in a MPC format, we obtain an efficient method for accommodating both the fundamental controller objective as well as run-time performance considerations.

Other mathematical developments [28, 29, 32, 33, 53, 54] of the satisficing concept are motivated by the desire to make robust decisions in the presence of uncertainty. These developments compare a utility defined over the consequences of a decision to a decision threshold. This decision threshold depends only on nature

[5] Note that because of our sign convention on v_R, time to contact is positive when relative velocity is positive.

and not on decision consequences. SDT is similar to these other developments in that it addresses robustness but, by contrast, SDT compares two utilities defined over the consequences of a decision whence SDT mathematically generalizes these decision rules (i.e., the decision threshold depends upon both control actions and the state of nature). However, these developments are more mature than SDT and include an axiomatic treatment of satisficing [29] which has not been generalized in SDT. It appears that an axiomatic treatment for SDT can be obtained by modifying the axioms in [29] using concepts from multi-attribute utility theory.

In addition to the abstraction hierarchy of human behavior developed by Rasmussen [8, 38, 40, 41, 46], the field of intelligent control relies on multi-resolutional behavior generation [1, 2, 23, 34, 44]. The foundational idea behind these developments is that a hierarchy of multi-resolutional modules is an efficient way to encode knowledge, to interpret sensory information, and to generate purposeful behavior. We employ this multi-resolutional encoding, and focus on the communication between levels in the hierarchy.

We model automobile driver longitudinal-control behavior by generating dedicated controllers for each necessary skill, and then manage these behaviors via a higher level controller. This model recalls work in behavior-based robotics wherein multiple low-level behaviors are fused and/or modulated by higher level controllers [5–7, 10, 12, 27]. Currently, our model employs deliberate switching between low-level skill-based controllers. An important area of future research in the spirit of behavior-based robotics, is the concurrent execution and behavior-based coordination of multiple skills.

1.4 Paper Outline

The paper is outlined as follows. Section 2 describes the multiple mental model framework, discusses SDT and MPC, and thereby provides a theoretical context to motivate and interpret experimental results. Section 3 identifies plausible perceptual cues and characterizes acceptable braking dynamics using measurements from both automated and manual responses to a vehicle cut-in event. Evidence is presented that supports T_c^{-1} and T_h as not only the key RB decision variables, but also the SB control variables that determine acceptable braking dynamics. Furthermore, evidence is presented to indicate that acceptable longitudinal braking is characterized by smooth, predictable trajectories in perceptual phase space. Section 4 presents experimental results from a simulator study that validate RB task management via SDT with identifiable and interpretable parameters. Section 5 presents professional driver responses to experimental cut-in events, identifies MPC parameters for emulating SB task execution, and discusses limitations to the MPC approach. The remaining section presents some suggestions for further work, a summary of the experimentally-supported theory, and conclusions.

2 Theory

2.1 Model Overview

We have been developing and continue to develop a suite of perception-based closed-loop models to emulate various SB driving behaviors (such as time-to-tangent-point curve negotiation [3], MPC-based braking [16], and time-to-lane-crossing-based lane-keeping [4]). To predict and describe driver behavior, it is useful to identify computational mechanisms for coordinating a set of such SB behaviors. One important aspect of this coordination is a method that predicts when a driver switches between different SB agents (i.e., how behavior is determined)[6]. For example, we are interested in conditions that trigger a switch from speed regulation to collision avoidance behaviors and, in the future, in those conditions when attention can be switched from longitudinal control to car phone usage (see Figure 4). In this section, we present a framework for top-down control of and communication between mental models including inputs and outputs for each level in the hierarchy. A preliminary discussion of how attention can be driven by bottom-up controller characteristics and how the perceptual and environmental bandwidths are related is provided in Appendix A. Within this paper, we focus primarily on how behavior may be generated within the hierarchical model framework. For each agent in the hierarchy, sensory information is interpreted and behavior is generated via the mental model $\mathcal{M}$. Notationally, state, behavior, and consequence for agents in each level of the hierarchy are represented by subscripts KB, RB, and SB for knowledge-based, rule-based, and skill-based, respectively. For example, a mental model agent at the rule base level has state θ_{RB}, control u_{RB}, and consequences c_{RB}.

[6]A second important aspect is how attention is shared between agents (i.e., how perception is controlled). Attentional sharing is necessary because drivers have limited computational and memory resources. A simple attentional model schedules attention between agents. More realistic models for attention are an area of future research.

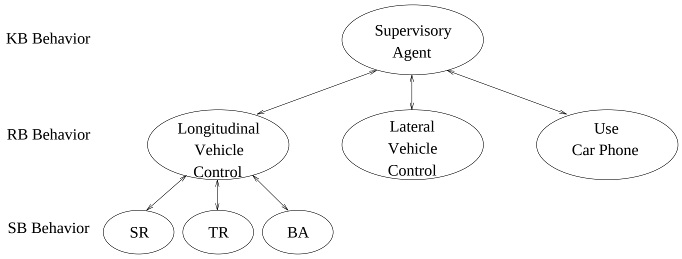

Figure 4. Hierarchical structure of agents in mental model society.

In describing agents at each level, we describe the agents role in the agent society, identify its inputs and outputs, and discuss methods for computationally modeling agent decision-making. Additionally, for RB and SB agents, we discuss the internal state used by the agent and present factors that determine acceptable agent behavior. For ease of reading, we underline these factors in the subsequent discussion.

2.1.1 KB Mental Models

The top nodes in the hierarchy are KB agents which organize and direct the agent society. For example (see Figure 4), in driving there is a need to coordinate between multiple tasks such as steering while talking on the car phone. The role of a KB agent includes managing attention and coordinating RB agents. Inputs to the KB agent include θ_{RB} and attentional requirements from RB agents. Outputs from the KB agent include the u_{KB} command to enable an RB agent and an attentional allotment to the RB agent. Computational modeling of KB agents, including representing goals and preferences, generating u_{KB} and θ_{KB}, and scheduling attention, is an area of future research (see [8, 38] for ideas).

2.1.2 RB Mental Models

The intermediate nodes in the hierarchy are RB task managers which coordinate multiple skills by switching. For example (see Figure 4), driving tasks that must be managed include longitudinal control, lateral control, and using the car phone. The role of an RB agent is to determine which SB controller to enable, when to switch from one SB agent to another, and which sensors should be consulted to reduce uncertainty and ensure satisficing performance (defined below). Inputs to RB agents include θ_{SB}, attentional requirements from the SB agent, and an attentional allotment from the KB agent. Outputs from the RB agents include the u_{RB} command to enable an SB agent, an attentional allotment to the SB agent, and θ_{RB} to the KB agent. Computational modeling of RB agents is performed using SDT [18] whereby we can partition the perceptual state space into regions; for each region, at least one SB controller is appropriate.

The state of the environment $\theta_{\mathrm{RB}} = \{\boldsymbol{\chi}, \boldsymbol{\psi}\}$ is used for monitoring SB behavior, and consists of two elements: (a) a perceptual state $\boldsymbol{\chi}$ used by the enabled SB controller to execute the assigned task ($\boldsymbol{\chi} = \theta_{\mathrm{SB}}$), and (b) perceptual cues $\boldsymbol{\psi}$ from disabled but engaged SB agents which are used to facilitate switches between SB behaviors[7]. To disable one SB agent and enable another SB agent, the RB agent must identify when the currently enabled SB agent cannot accomplish the assigned RB task. Acceptable RB agent behavior is characterized by flexible and appropriate SB agent selection. Since switching between SB agents requires both attention and effort, switching should be carefully managed. Flexible agent selection not only requires a tolerance for actual and expected SB agent behavior, but also well-defined conditions under which an SB agent should be disabled.

2.1.3 SB Mental Models

The terminal nodes in the hierarchy are SB controllers which execute the task specified by the RB agent. For example (see Figure 4), in longitudinal control there include three closed loop controllers: Speed Regulation (SR) wherein the driver regulates speed about a desired value, Time headway Regulation (TR) wherein the driver follows another vehicle at a desired time headway, and Brake to Avoid collision (BA)

[7] The difference between $\boldsymbol{\chi}$ and $\boldsymbol{\psi}$ is that estimates of $\boldsymbol{\chi}$ can be obtained the internal model used to generate u_{SB}, whereas estimates of $\boldsymbol{\psi}$ must be actively sensed (via visual angle or rate of optical expansion) and cannot be predicted in open loop (through internal model of time to collision dynamics).

wherein the driver reacts to significant dynamic disturbances such as emergency braking by a lead vehicle[8]. The role of an SB agent is to execute a perception-based control law that accomplishes the performance objective. Inputs to SB agents include sensory observations of the environment and an attentional allotment from the RB agent. Outputs from SB agents include the u_{SB} behavior command, an attentional requirement to the RB agent, and θ_{SB} to the RB agent. MPC (a variant of the optimal control models successfully employed to emulate skill-based linear control [46]) emulates skill-based nonlinear control and can be used in computational modeling of SB agents.

The state of the environment for SB agents is a dynamic percept sufficiently informative (i.e., stabilizable and detectable [26]) for closed loop control. This closed loop control requires an observation of current state and an estimation of future states as a function of action. The perceptual state must therefore be perceptually feasible (observable through existing senses) and cognitively plausible (usable in perception-response control). Acceptable SB agent behavior is predictable by other drivers, efficient in attentional usage, and reliable in accomplishing the assigned task.

2.2 RB Task Management via SDT-based Task Switching

Many cognitive scientists recognize that insistence on optimality is a misplaced requirement in situations of limited resources and information, and that optimality inadequately describes observed behavior in naturalistic settings [14, 48, 49, 57]. Simon [47] addressed the issue of limited or bounded rationality by defining an aspiration level, such that once this level is met, the corresponding solution is deemed adequate, or *satisficing*. An important characteristic of Simon's satisficing principle is that decisions are deemed adequate on the basis of a *comparison*: any decision which exceeds the aspiration level is admissible. We employ this comparative characteristic by constructing and comparing two set membership functions similar to the way benefit and cost are compared in economics literature. The key to this development lies in partitioning preferences over consequences into a generalized type of benefit called *accuracy*, and a generalized type of cost called *rejectability*. These two decision attributes may be operationally characterized as follows:

ACCURACY: A natural characterization of the benefit of a decision is accuracy, meaning *conformity to a standard.* In practical contexts, the standard corresponds to whatever goal or objective is relevant to the problem, and accuracy corresponds to the degree of success in achieving that goal. When represented using fuzzy logic, the term *accuracy* refers to the set membership function associated with the linguistic variable ACCURATE[9].

REJECTABILITY: Actions may also be evaluated strictly in terms of their undesirable consequences or rejectability[10], meaning *susceptibility or exposure to something undesirable.* Typically, these consequences may be manifest in the form of costs or other penalties that would accrue simply as a result of taking action, regardless of its accuracy. Rejectability corresponds to the degree to which actions accrue costs or penalties. When represented using fuzzy logic, the term *rejectability* refers to the set membership function associated with the linguistic variable REJECTABLE.

For example, in regulator design, the fundamental objective is to drive the system to and maintain the system at a desired operating point. Thus, accuracy refers to the degree to which the possible controlled states satisfy this objective. Independent of the desire to regulate the system is the desire to prevent excessive control authority and oscillatory state transitions. Thus, rejectability refers to the cost of possible controlled states with respect to these undesirable consequences.

Using Levi's error avoidance principle [25], SDT provides a method by which the accuracy and rejectability set membership functions can be merged: *to avoid error, a decision maker accepts those decisions which are* ACCURATE *and* not REJECTABLE. Formally, let U denote the set of possible decisions or actions, and let Θ denote the states of nature. For each decision $u \in U$ and for each state of nature $\theta \in \Theta$, a consequence results which is the effect of making decision u when nature is in state θ. The accuracy $\mu_A : U \times \Theta \mapsto \Re$ and rejectability $\mu_R : U \times \Theta \mapsto \Re$ set membership functions are preference relations defined for each consequence (i.e., action/state-of-nature pair).

[8]We do not consider alternative collision avoidance strategies such as swerving because these strategies emerge from the interaction between multiple RB agents. Instead, we leave these areas for future work on the fusion of RB agent behaviors.

[9]For the remainder of the paper, we use capital letters and a separate font when we refer to linguistic variables, but will make no such distinction for membership functions.

[10]In [18, 51], the term rejectability is used by the more colloquial term liability.

In SDT, the set of all decisions which cannot be justifiably eliminated is called the *satisficing set*, and is defined as (see [17])

$$\mathsf{SATISFICING} = \mathsf{ACCURATE} \text{ and } \mathsf{not}(\mathsf{REJECTABLE}).$$

For the problems addressed herein, we wish to include multiplicative hedges $\alpha, \rho \in [0,1]$ which allow the fuzziness inherent in the consequences of an action to be parameterized. Thus, we form the satisficing membership function as

$$\mu_{S_b} = \alpha\mu_A \star (1 - \rho\mu_R) \tag{2}$$

where b represents a design parameter which is related to ρ and α, where $\star$ represents a t-norm[11], and where $1-\rho\mu_R$ represents the complement of the hedged set $\rho\mu_R$. When $\star$ represents the t-norm $Y \star Z = \max(0, Y+Z-1)$ (see, for example, [11]), the satisficing set membership becomes

$$\begin{aligned} \mu_{S_b} &= \max(0, \alpha\mu_A - \rho\mu_R) \\ &= \max(0, \mu_A - b\mu_R), \end{aligned} \tag{3}$$

where $b = \rho/\alpha \in [0, \infty)$ is called the *rejectivity* and parameterizes the relative weight[12] between accuracy and rejectability. The comparative nature of (3) is best illustrated by considering the region of support (area of nonzero set membership) for the satisficing set, which is given by

$$S_b = \{(u;\theta) : \mu_A(u;\theta) \geq b\mu_R(u;\theta)\}. \tag{4}$$

As discussed in Sections 1 and 2, we employ satisficing decision theory (SDT) to encode the role of RB agents. Recall that a mental model $\mathcal{M}$ consists of a set of controls u, a set of perceptual states θ, and an ordered set of consequences $c = (u, \theta)$. In SDT, preferences over consequences are represented by the benefit-like *accuracy* attribute and the cost-like *rejectability* attribute. These attributes are compared to determine when action u is admissible given state θ (i.e., when consequences are satisficing). Formally, the set of RB actions U_{RB} consists of an enabling command to one and only one of the SB agents whence, for the task of longitudinal control shown in Figure 4, $U_{\mathrm{RB}} = \{\mathrm{TR}, \mathrm{SR}, \mathrm{BA}\}$. Given the set of perceptual states Θ_{RB}, the accuracy function $\mu_A : U_{\mathrm{RB}} \times \Theta_{\mathrm{RB}} \mapsto \Re$ and the rejectability function $\mu_R : U_{\mathrm{RB}} \times \Theta_{\mathrm{RB}} \mapsto \Re$ are compared to determine the set of satisficing consequences [18]

$$S_b = \{(u_{\mathrm{RB}}, \theta_{\mathrm{RB}}) : \mu_A(u_{\mathrm{RB}}, \theta_{\mathrm{RB}}) \geq b\mu_R(u_{\mathrm{RB}}, \theta_{\mathrm{RB}})\}. \tag{5}$$

Given (5), we can restrict attention to those states of nature which are satisficing for a given u_{RB}, and those controls which are satisficing given the state of nature, respectively defined as

$$\begin{aligned} S_b(u_{\mathrm{RB}}) &= \{\theta_{\mathrm{RB}} : \mu_A(u_{\mathrm{RB}}, \theta_{\mathrm{RB}}) \geq b\mu_R(u_{\mathrm{RB}}, \theta_{\mathrm{RB}})\} \\ S_b(\theta_{\mathrm{RB}}) &= \{u_{\mathrm{RB}} : \mu_A(u_{\mathrm{RB}}, \theta_{\mathrm{RB}}) \geq b\mu_R(u_{\mathrm{RB}}, \theta_{\mathrm{RB}})\}. \end{aligned}$$

In terms of task management by a RB agent, suppose a SB agent $\alpha \in U_{\mathrm{RB}}$ is enabled. The RB agent monitors θ_{RB}, and when $\theta_{\mathrm{RB}} \in S_b(\alpha)$ no change is necessary (although a KB agent can dictate a change). However, when $\theta_{\mathrm{RB}} \notin S_b(\alpha)$, the current SB controller is not acceptable and must be switched to a controller that is appropriate for the circumstances. Given the need to switch, any $u_{\mathrm{SB}} \in S_b(\theta_{\mathrm{RB}})$ can be employed[13]. An algorithm can be outlined for such task management as follows:

If $\theta_{\mathrm{RB}} \in S_b(u_{\mathrm{RB}})$

$u'_{\mathrm{RB}} = u_{\mathrm{RB}}$

Else

$u'_{\mathrm{RB}} \in S_B(\theta_{\mathrm{RB}})$

[11] A t-norm is the fuzzy instantiation of a logical "and" conjunction [11, 31].

[12] The subjective selection of this relative weighting is analogous to the tradeoff between the size and power of a statistical hypothesis test using Neyman-Pearson decision theory. Similar to the way in which subjectively selecting a test's size determines the test's power in Neyman-Pearson hypothesis testing, subjectively selecting b determines the relative importance of accuracy and rejectability.

[13] Note that because several SB agents can be satisficing for a given state θ_{RB}, chattering between controllers is avoided and, instead, replaced by the hysteresis effect noted in [4].

2.3 MPC-Based SB Behavior

Skill-based behavior can be accomplished by generating a control u_{SB} using a perception-based response. For longitudinal control, braking force and throttle must be controlled whence, from Figure 3, $u_{\mathrm{SB}} = u_A$. In the absence of disturbances, when the driver for vehicle A executes control u_A given a current perceptual state $\theta_{\mathrm{SB}} = \boldsymbol{\chi}(k)$, a new perceptual state

$$\boldsymbol{\chi}(k+1) = \boldsymbol{f}(\boldsymbol{\chi}(k), u_A, u_B) \tag{6}$$

results. MPC methods, which do not appear to have been extensively studied in modeling driver behaviors, are compatible with many of the criteria (such as the ability to handle constraints, employ a finite prediction horizon, and track a time-varying virtual target) necessary for modeling perceptual regulation behaviors. Given a desired perceptual state $\boldsymbol{\chi}^*$, the model predictive controller [30, 39, 45] is obtained by minimizing the cost function

$$J_N = \phi(\boldsymbol{\chi}(N), \boldsymbol{\chi}^*) + \sum_{k=0}^{N-1} L(\boldsymbol{\chi}(k), \boldsymbol{\chi}^*, u_A(k)) \tag{7}$$

with respect to the control sequence $u_A(0), \ldots, u_A(N-1)$ subject to the control bounds given by

$$u(k) \in U_A \text{ for all } k \qquad \Delta u(0) \in U_A^{\Delta}, \tag{8}$$

and the dynamics constraint given by (6). In (7), the function ϕ represents the penalty for terminating control at time N in state $\boldsymbol{\chi}(N)$, and L represents the incremental penalty for being in state $\boldsymbol{\chi}(k)$ and using $u_A(k)$ at time k. From the sequence obtained by minimizing (7), the first control $u_A(0)$ is injected into the plant $\boldsymbol{f}$, and the constrained minimization is repeated for the next (and all future) time step(s).

Recall that our definition of a mental model includes state θ, action u, and ordered consequences c. For SB longitudinal control, the action $u_{\mathrm{SB}} = u_A$, the state θ_{SB} includes the current perceptual state $\boldsymbol{\chi}$, the consequences are future perceptual states $\boldsymbol{\chi}(k+\ell)$, and the preference ordering is performed by the cost function J_N.

For example, in longitudinal control we will use the RB state $\theta_{\mathrm{RB}} = [T_c^{-1}, T_h, v_A]^T$. Depending on which SB agent is enabled, one of several SB perceptual states $\boldsymbol{\chi}$ are possible. In the absence of other traffic, the driver regulates vehicle speed v_A (i.e., the enabled SB agent is SR) around an operating point v_A^* subject to the satisficing constraint set whence the desired perceptual state is $\boldsymbol{\chi}^* = v_A^*$, the perceptual state is $\boldsymbol{\chi} = \theta_{\mathrm{SB}} = v_A$, and engaged but disabled states are $\boldsymbol{\psi} = [T_c^{-1}, T_h]^T$. In the presence of other traffic, the driver avoids collision or regulates time headway (i.e., the enabled SB agent is either TR or BA) whence the desired perceptual state is $\boldsymbol{\chi}^* = [0, T_h^*, \cdot]^T$, the perceptual state is $\theta_{\mathrm{SB}} = \boldsymbol{\chi} = [T_c^{-1}, T_h]^T$, and engaged but disabled states are $\boldsymbol{\psi} = v_A$.

For longitudinal control, the following factors influence the amount of braking force that can and will be applied:

- Behavior of lead vehicle u_B.
- Relative importance of making $T_c^{-1} \to 0$ versus making $T_h \to T_h^*$.
- Cost of control ("feel" of deceleration and effort to change pedal position(s)).

These factors can be used in the MPC formalism by minimizing the cost function (7) with

$$\begin{aligned} \phi(\boldsymbol{\chi}(N), \boldsymbol{\chi}^*) &= [\boldsymbol{\chi}(N) - \boldsymbol{\chi}^*]^T P [\boldsymbol{\chi}(N) - \boldsymbol{\chi}^*] && (9) \\ L(\boldsymbol{\chi}(k), \boldsymbol{\chi}^*, u_A(k)) &= [\boldsymbol{\chi}(k) - \boldsymbol{\chi}^*]^T Q [\boldsymbol{\chi}(k) - \boldsymbol{\chi}^*] + h(u_A(k))^T S h(u_A(k)) + \\ &\quad \Delta u_A(k)^T S_\Delta \Delta u_A(k) && (10) \end{aligned}$$

where $u_A(k)$ represents the pedal positions, and $h(u_A(k))$ represents the acceleration generated by the pedal position $u_A(k)$. The elements of Q and P balance the relative importance of time headway and time-to-collision, u_B influences behavior via the equation governing $\boldsymbol{\chi}(k)$'s dynamics (see (6), and S and S_Δ govern the "feel" of the deceleration and acceleration dynamics (S penalizes controls that produce large accelerations, and S_Δ penalizes controls that require the driver to drastically change foot position).

3 Experiment 1: Longitudinal Control Test Results

The key points made in this section are as follows:

- The perceptual sub-state $\chi = [T_c^{-1}, T_h]^T$ can be used to efficiently classify behavior into braking and nominal categories.
- Drivers appear to first establish infinite time to collision ($T_c^{-1} \leq 0$) and then drive the system to a desired time headway T_h^*. Acceptable behaviors are characterized by a counterclockwise perceptual phase trajectory, and unacceptable trajectories by either a clockwise or a non-smooth trajectory.

3.1 Experiment Description

In the experiment, two vehicles, denoted vehicle A and vehicle B as in Figure 3, drive on a test track in adjacent lanes. Vehicle B passes vehicle A, slows down, and then cuts into vehicle A's lane at a time that is unknown to vehicle A. During vehicle A's reaction to this cut-in event, vehicle A records its velocity v_A, its brake pressure (no throttle measurements were recorded), the range R measured with a three beam laser radar, and relative velocity v_R obtained by processing the three range measurements. The experiment was conducted using both an automated longitudinal control system as well as a professional driver responding to cut-in events. The data is grouped into three categories: acceptable automated performance, unacceptable automated performance, and professional driver performance without automation, as described in Table 2.

DATA FILES	DESCRIPTION
`auto_0.prn` - `auto_9.prn`	acceptable automated performance
`auto_ng.prn`	unacceptable automated performance
`man_0.prn` - `man_3.prn`	professional driver performance without automation

Table 2. Description of preliminary data sets. Sample time is 10 ms.

3.2 Observations

It is helpful to illustrate the time histories and perceptual phase plane trajectories (see Appendix B for a tutorial on phase plane trajectories) subsequent to the cut-in event for each data class identified in Table 2. Figures 5-10 display the perceptual trajectory using the sub-state $\chi = [T_c^{-1}, T_h]^T$ for three representative trials. In the perceptual phase plane figures, the trajectories are shown only after a cut-in event (detected by observing a discontinuity in R); the large diamond indicates the initial perceptual state that results from the cut-in.

The data is classified into two categories: those for which active braking occurs and those for which no such braking occurs, indicated in the Figures with a $\times$ and a $\circ$, respectively. The sequence of $\circ$'s present after the cut-in event indicate the amount of time taken to react to the cut-in event.

Three observations are apparent:

- Braking is initiated (ignoring reaction) when $T_c^{-1} > 0$ ($v_R > 0$). Conversely, a driver is likely to accelerate when $T_c^{-1} < 0$ ($v_R < 0$). Thus, dividing driver behavior into active braking and nominal (not-active) braking produces a division roughly at $T_C^{-1} = 0$ ($v_R = 0$).
- When $T_c^{-1} \leq 0$, the factor determining dynamic driver behavior appears to be related to time headway. This is observable from the driver response in `man_0` wherein the driver first establishes zero relative velocity and then appears to regulate vehicle speed around the time headway value $T_h^* \approx 1.65$s.
- The characteristics of the phase plane trajectory influence the acceptability of the automated performance. Each trajectory in the set of *acceptable* automated behaviors (`auto_0`-`auto_9`) and the manual behaviors (`man_0`-`man_3`) exhibit a *counterclockwise* movement in the phase plane, but the *unacceptable* automated behavior `auto_ng` exhibits a *clockwise* movement. This is most evident when `auto_ng` is compared to `auto_6`, since these two have similar initial conditions. The test driver reported that `auto_ng` was unacceptable because the braking action was too extreme given that the relative velocity,

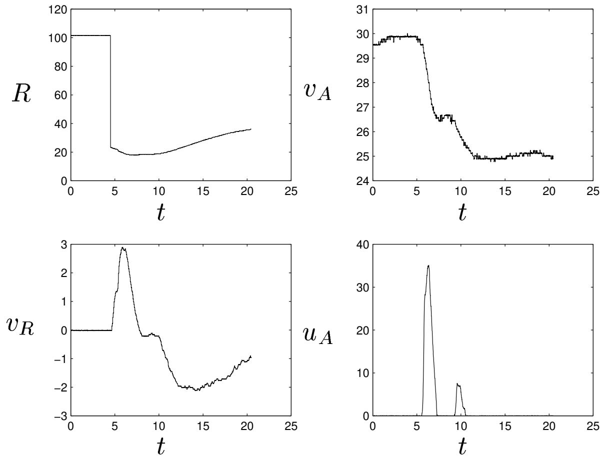

Figure 5. Time histories of `auto_6.prn`.

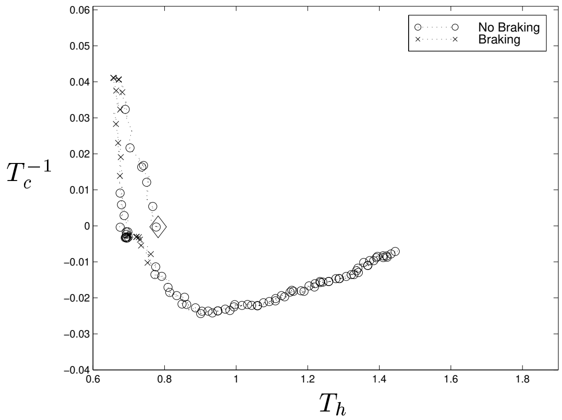

Figure 6. Perceptual phase trajectory of `auto_6.prn`.

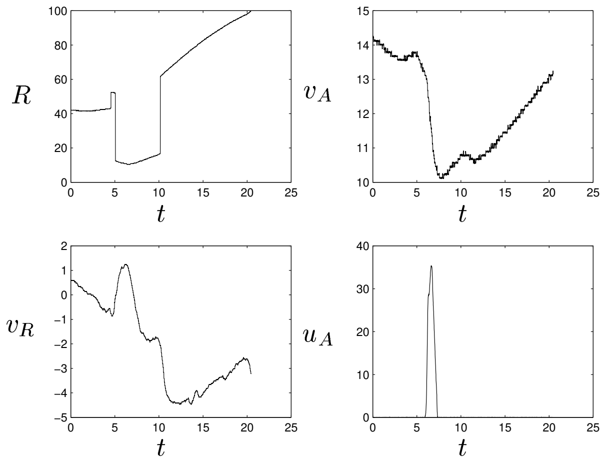

Figure 7. Time histories of `auto_ng.prn`.

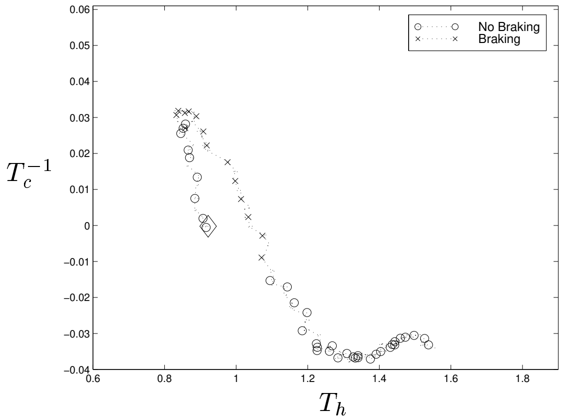

Figure 8. Perceptual phase trajectory of `auto_ng.prn`.

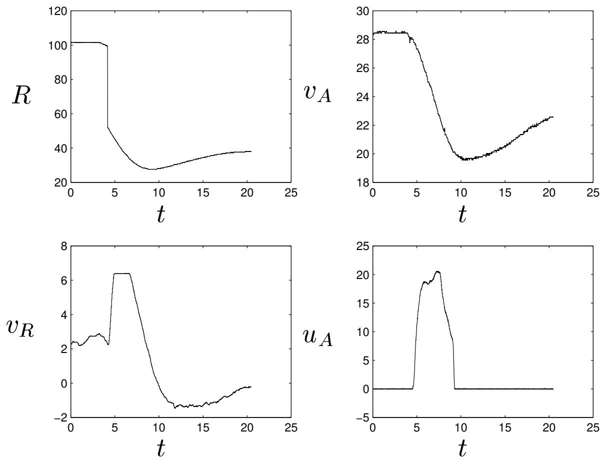

Figure 9. Time histories of `man_0.prn`.

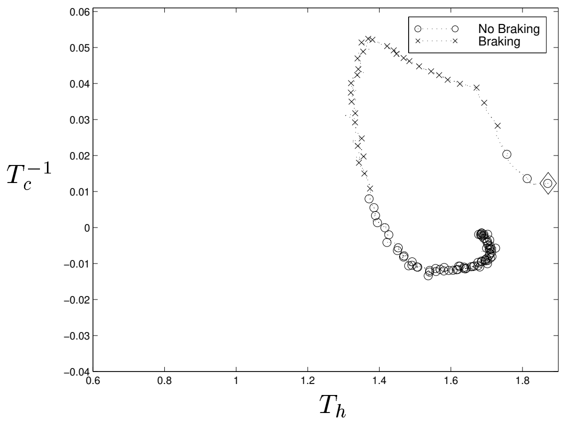

Figure 10. Perceptual phase trajectory of `man_0.prn`.

though positive, was small in magnitude. This "hard-braking plus low relative velocity" characteristic is manifest as a clockwise trajectory in the perceptual phase plane.

These observations support the hypothesis that drivers employ T_c^{-1} and T_h to determine braking response. The third perceptual state v_A is important for speed regulation, such as driving in low traffic density.

3.3 Results: Perceptual States and Brake Initiation

Two important empirical goals are (1) to motivate a valid perceptual state and (2) to determine conditions when braking is initiated. We address these issues by first identifying feasible perceptual cues from the literature, second using a correlation analysis, and finally using a nonlinear regression analysis.

3.3.1 Perceptual States

Since T_c^{-1} can be detected by humans (consider the Weber ratio model for looming detection), this perceptual state is perceptually feasible [24, 56]. Similarly, T_h can be described as a time to collision with road textures and is thus perceptually feasible. Finally, v_A is perceptually feasible since the driver can learn to associate speedometer readings with optic flow of road and world textures. Thus, we hypothesize that $\theta_{\mathrm{RB}} = [T_c^{-1}, T_h, v_A]^T$, and gather evidence to support this hypothesis including evidence that the subspace $\mathrm{sp}[T_c^{-1}, T_h]^T$ is important to determine when to brake and with what dynamics.

3.3.2 Correlation Analysis

One approach to motivating a perceptual state is determining first order relationships between braking and sensory inputs. Support for such first order relationships can be obtained by a correlation analysis of variables with brake pressure u_A. Since measurements include braking force (plus noise) but not throttle position, an analysis of correlation coefficients will not provide a conclusive motivation for a particular perceptual state. Nevertheless, such an analysis can help identify candidate variables. The correlation coefficients obtained from time series data subsequent to a cut-in event are tabulated in Table 3. From this

Data Set	T_h	v_R	v_A	R	T_c^{-1}
`man_0`	-0.64	0.90	0.54	-0.02	0.95
`man_1`	-0.48	0.82	0.17	-0.23	0.88
`man_2`	0.29	0.94	0.64	0.44	0.96
`man_3`	-0.06	0.94	0.80	0.45	0.89
combined	-0.22	0.90	0.80	0.45	0.92

Table 3. Correlation coefficients with braking force after cut-in events.

table, we identify v_R and T_c^{-1} as those variables most highly correlated with braking force. It also appears that v_A is highly correlated with braking force, but that T_h has little effect. However, if we look at correlation coefficients under all operating conditions (not just subsequent to cut-in events), we see a deemphasis on v_A and an increased emphasis on T_h (although not convincing evidence).

Data Set	T_h	v_R	v_A	R	T_c^{-1}
`man_0`	-0.35	0.77	0.12	-0.26	0.91
`man_1`	-0.48	0.82	0.17	-0.17	0.88
`man_2`	-0.29	0.91	0.09	-0.27	0.93
`man_3`	-0.21	0.94	0.26	-0.13	0.89
means	-0.33	0.86	0.16	-0.22	0.90

Table 4. Correlation coefficients with braking force under all operating conditions.

The importance of T_h is most convincingly illustrated by looking at the (T_h, T_c^{-1}) phase plane plots for the manual data. The effect, typified in Figure 10, shows that regulation is likely done about some $T_c^{-1} = 0$ ($v_R = 0$) and $T_h \approx 1.65$ operating point. Thus, we find evidence that $[T_c^{-1}, T_h]^T$ (T_c^{-1} for brake initiation, see the correlation coefficients in Tables 3-4, and T_h for regulation, see the dynamic response in Figure 10 and note that other manual responses exhibit similar characteristics) is a relevant set of measurements for determining when drivers brake. We provide additional evidence that T_h is important in Sections 4 and 5.2.

3.3.3 Regression Analysis

To identify the driver response to a cut-in event, we look for a subset of $\theta_{\mathrm{RB}} = [T_c^{-1}, T_h, v_A]^T$ that allows us to distinguish between active braking and nominal (speed or time headway regulation) behaviors. To accomplish this, we construct a regression model to identify variables relevant to braking force. This regression model is then used to classify perceptual states into active braking and nominal classes. We suggest that active braking indicates a perceptual state that is not satisficing for nominal behavior (for speed and time headway regulation $\theta_{\mathrm{RB}} \notin S_b(u_{\mathrm{RB}}),\ u_{\mathrm{RB}} \in \{\mathrm{SR}, \mathrm{TR}\}$).

From the sign convention introduced in the introduction, the control variable u_A is negative when the driver/automation is actively braking, and non-negative when no such braking occurs. Since the data measurements include brake pressure but do not include any information regarding accelerator pedal (such as throttle opening), u_A cannot be larger than zero. Since the brake pressure data is noisy, we employ the empirically motivated threshold of 0.1 kg/cm^2 to classify measured behavior, whence

$$u_A \in \begin{cases} \text{Nominal} & u_A \le 0.1 \\ \text{Braking} & u_A > 0.1 \end{cases}.$$

The threshold of 0.1 kg/cm^2 is selected because a measurement noise with three sigma limit of approximately ± 0.1 is added to the measured braking forces; i.e., $u_A^{\text{observed}} = u_A^{\text{actual}} + \eta$, η zero mean with variance $(0.1/3)^2$. This threshold effectively prevents noise from biasing the classification results.

Due to the absence of accelerator pedal information, we employ a saturating linear model in the regression analysis

$$\hat{u}_A = \min\{0, \phi^T \boldsymbol{\chi}\} \tag{11}$$

where $\boldsymbol{\chi}$ is the perceptual state, and where ϕ is the vector of regression parameters. The regressors $\hat{\phi}$ are the best least squares fit obtained using MATLAB's `nlinfit` function. Given a perceptual state $\boldsymbol{\chi}$ and a vector of regressors $\hat{\phi}$, the model predictions are classified as

$$\hat{u}_A \in \begin{cases} \text{Nominal} & \hat{\phi}^T \boldsymbol{\chi} \ge 0 \\ \text{Braking} & \hat{\phi}^T \boldsymbol{\chi} < 0 \end{cases}$$

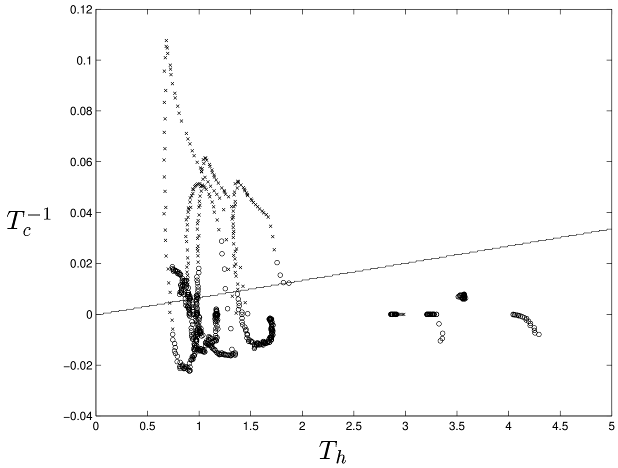

Figure 11. Decision planes obtained via regression for $\boldsymbol{\chi} = [T_c^{-1}, T_h]$.

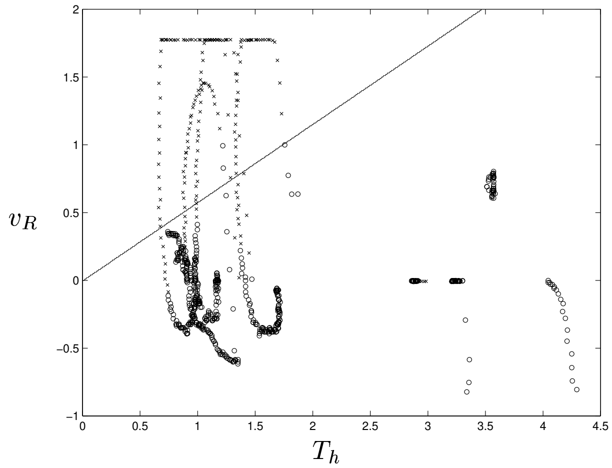

Figure 12. Decision planes obtained via regression for $\boldsymbol{\chi} = [T_h, v_R]$.

In Table 5, we report classification[14] results (diagrammed in Figures 11-12 with $\times$ indicating braking, $\circ$ indicating nominal, and the solid line indicating the decision plane) for $\boldsymbol{\chi} = [T_h, v_R]^T$ and $\boldsymbol{\chi} = [T_c^{-1}, T_h]^T$.

[14] It is possible to use any of the variables in Table 3-4 as regressors. We present result for two of the most successful pairs used in classification and omit less successful variables.

It should be noted that the higher classification accuracy for $\chi = [T_h, v_R]^T$ does not indicate that these

χ	Data Set	% Misclassified	% False Brake	% Missed Brake
$[T_h, T_c^{-1}]^T$	`man_0`	1.46	1.22	0.24
	`man_1`	4.25	4.25	0.00
	`man_2`	19.43	18.80	0.63
	`man_3`	22.61	16.50	6.10
$[T_h, v_R]^T$	`man_0`	4.64	4.64	0.00
	`man_1`	4.93	3.52	1.42
	`man_2`	2.98	2.98	0.00
	`man_3`	8.54	8.54	0.00

Table 5. Classification accuracy obtained via regression for field test data.

perceptual variables are superior. To obtain enough sample points, the classification is performed over the entire braking trajectory, even though the objective of the classification is determine the regions when braking is initiated. Thus, several braking conditions are a result of the dynamical response of the driver rather than of the unsatisfactory nature of the perceptual state. (Additionally, recall that T_c^{-1} appears to be directly perceived and is thus perceptually feasible.) We will return to these issues when we report the experiments from the CBR simulator, but for now we emphasize that there is evidence which supports the hypothesis that a curve in sp$[T_c^{-1}, T_h]$ distinguishes between braking and nominal behaviors.

3.4 Results: Dynamic Behaviors

In this section, we analyze the dynamic behavior of two automated controllers. Consider the dynamic

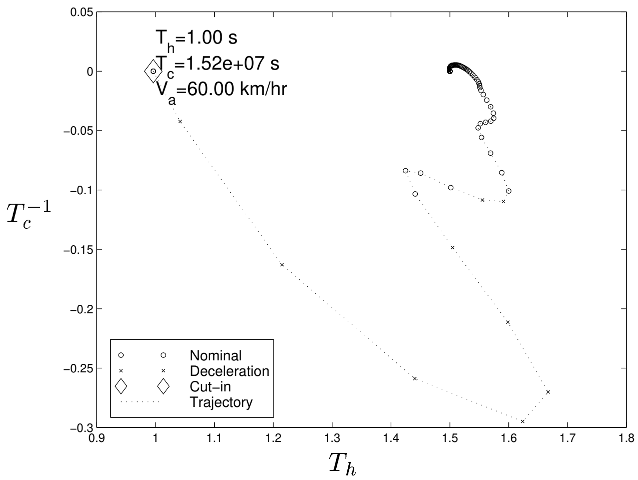

Figure 13. Phase trajectories of controller 1.

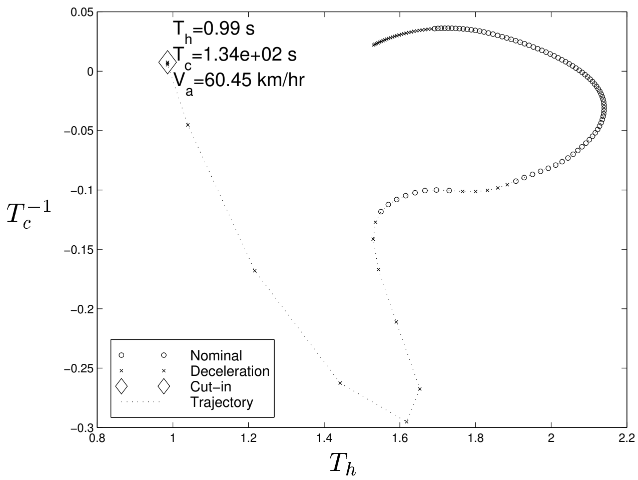

Figure 14. Phase trajectories of controller 2.

responses shown in Figures 13 and 14. These phase plane trajectories represent the automated dynamic response to a cut in event. Before the cut in, the automated controller has established a speed of $v_A \approx 60$ km/hr and a headway of $T_h^* \approx 1.5$ s. A vehicle with constant speed $v_B = 60$ km/hr cuts in at a headway of 1.0 s. For controller 1, we see that the desirable counterclockwise trajectory does not smoothly approach the desired time headway value, but instead has an undesirable bump. A bump in the perceptual trajectory corresponds to undesirable and unnatural deceleration. By contrast, for the same cut-in scenario controller 2's response is shown in Figure 14. The counterclockwise motion is exhibited, but this time the trajectory is more smooth and free of undesirable bumps. Thus, we predict that the latter automated control will seem more natural and acceptable to drivers than the former automated control.

Now consider the dynamic responses shown in Figures 15 and 16. The perceptual trajectories again represent the dynamic responses to a cut-in event. Before the cut-in occurs, the automated controller has

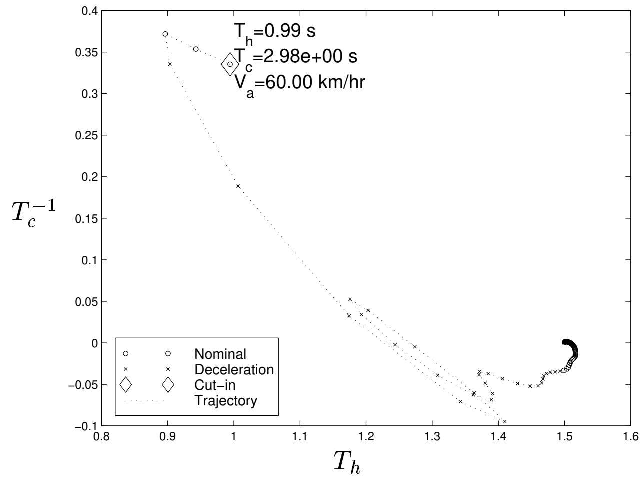

Figure 15. Phase trajectories of controller 1.

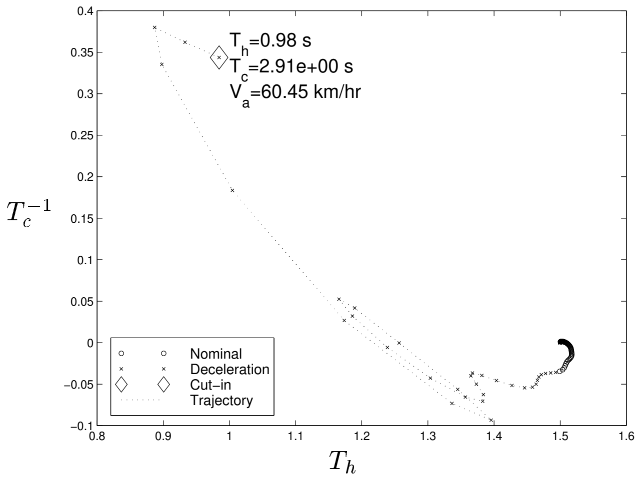

Figure 16. Phase trajectories of controller 2.

established a speed of $v_A \approx 60$ km/hr at a headway of $T_h^* \approx 1.5$ s. A vehicle traveling at the constant slower speed $v_B = 40$ km/hr cuts in at a headway of $T_h \approx 1.0$ s. For this case, both controller 1 and controller 2 have bumpy perceptual trajectories. Thus, the driver model predicts that the automated behavior for this and similar scenarios will seem uncomfortable and unnatural to the driver.

Potentially problematic cut-in scenarios are characterized by counterclockwise or non-smooth perceptual phase plane trajectories. Under conditions similar to those shown in Figures 15-16, subjective evaluations reported a surge in the vehicle behavior. This surge was caused by a braking interval followed by a coasting interval and then followed by another braking interval. This surge, though moderate, still "felt" unnatural. This subjective evidence indicates that perceptual phase plane trajectories should not only be counterclockwise, but should also be smooth.

4 Experiment 2: Simulator Experiment Results

The key points from this section are as follows:

- Smooth counterclockwise perceptual dynamics of braking are observed in the simulator studies.
- Physical interpretations of accuracy and rejectability membership functions are established and empirical estimates of these functions are obtained.
- Classification of braking and nominal behavior are successfully performed using SDT with the measured membership functions.

4.1 Experiment Description

The SIRCA simulated driving environment created by Marcos Fernandez from the University of Valencia in Spain includes approximately six miles of highway with three lanes in each direction and ambient traffic. In the experiment, a subject performs lateral control but engages a cruise control (CC) system to perform longitudinal control about a preset condition ($v^* \approx 20$m/s ≈ 43mph). During the experiment, a cut-in vehicle passes the subject's vehicle while the CC is engaged and cuts into the lane with a specified relative velocity $v_R(0)$ and fixed initial time headway $T_h(0)$ randomly selected from the experimental conditions $v_R(0) \in \{-10, -5, 0, 5\}$ (m/s) and $T_h(0) \in \{0.5, 1.25, 2\}$ (s). Subsequent to a cut-in event (after maintaining the desired cut-in speed for 10 seconds), the cut-in vehicle speeds away and disappears into the horizon. If the subject disengaged the CC in response to the cut-in, they restart the CC system and continue driving. Ten subjects, naive to the experimental purposes, participated in the experiment. T_h, v_A, v_B, lateral position, and steering were recorded, and data were partitioned into two classes: *active braking* (brake pedal depressed) and *nominal behavior* (CC engaged, accelerator depressed,or engine braking[15]).

[15]The subject must disengage the CC to implement engine braking.

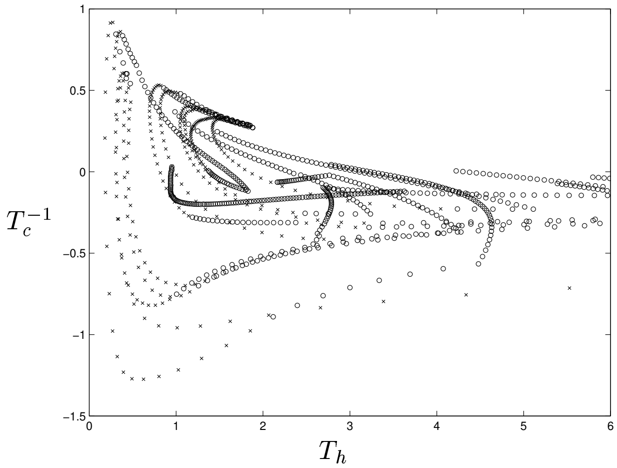

Figure 17. Perceptual phase plane trajectories for subject IT with initial conditions $(T_h(0), v_R(0)) = (2\text{s}, 10\text{m})$.

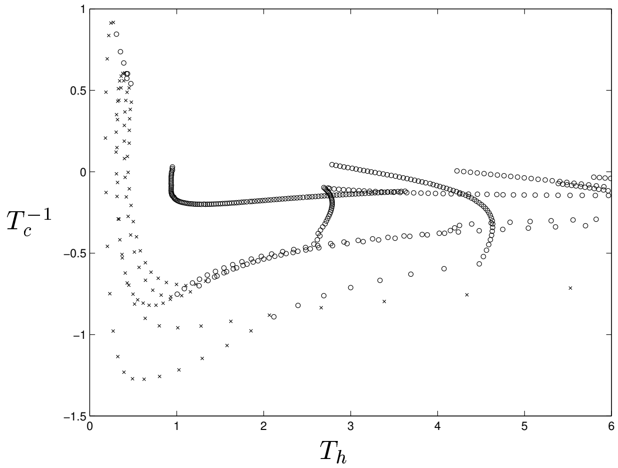

Figure 18. Perceptual phase plane trajectories for subject IT with initial conditions $(T_h(0), v_R(0)) = (0.5\text{s}, 5\text{m})$.

Responses for subject IT are shown in Figures 17-18 for $(T_h(0), v_R(0))$ given by, respectively, $(2\text{s}, 10\text{m})$ and $(0.5\text{s}, 5\text{m})$. The phase trajectory begins the instant the center of the lead vehicle crosses into the subject's lane. In the figures, $\times$ indicates braking and $\circ$ indicates nominal behavior. (A line of $\circ$'s moving toward the southwest corner of the plot occurs when IT manually turns the ACC system off, but does not brake. A line of $\circ$'s moving toward the northwest of the plot occurs when IT manually turns the ACC system back on and accelerates toward the lead vehicle).

4.2 Observations

Three observations are immediately apparent from the plots. First, drivers require a finite amount of time to react to the cut-in. For $(T_h(0), v_R(0)) = (2\text{s}, 10\text{m})$ the time to actively brake is between 0.5 and 2.3 seconds (subject IT frequently anticipated the cut-in event and manually switched off the CC to engage engine braking prior to actively pushing the brakes so the reaction times may not precisely represent the time taken to act). For $(T_h(0), v_R(0)) = (0.5\text{s}, 5\text{m})$ the time to actively brake is between 0.1 and 0.2 seconds (IT frequently applied the brakes before the center of the cut-in vehicle crossed into the lane whence the reaction times, measured as the difference between the time when braking is initiated and the time when center of vehicle B crosses into vehicle A's lane, appear artificially low). Second, subjects frequently brake much longer than is necessary to establish safe following distances. This probably results from "soft" simulator brakes and from the inability to sense braking force through physical cues. Third, braking patters exhibit the characteristic smooth counterclockwise movement identified in Section 3.

4.3 Empirically Derived Memberships

Recall from Sections 1-2 that longitudinal control is modeled by a cooperation between the RB task manager and a team of SB controllers. To determine how the longitudinal RB agent manages its assigned task, it is necessary to identify the actions and states of each agent involved in the cooperation. To identify possible agent actions, consider the set of SB controllers diagrammed in Figure 4. The longitudinal RB agent must enable one and only one of the SB controllers to execute control, whence u_{RB} is an enabling command to the SB agent such that $u_{\text{RB}} \in \{\text{TR}, \text{SR}, \text{BA}\}$. The corresponding SB agents control brake pressure and throttle opening whence $u_{\text{SB}} = u_A$, where u_A is the brake/accelerator control for vehicle A as diagrammed in Figure 3. To identify possible RB agent states, we use the perceptual state $\theta_{\text{RB}} = [T_c^{-1}, T_h, v_A]^T$ identified in Section 3[16].

The key to understanding the concepts of accuracy (μ_A) and rejectability (μ_R) is found in the notion of a utility. Loosely speaking, a utility is a numerical representation of a person's subjective values (i.e., subjective values are the embodiment of the preferences among consequences). The accuracy membership function is a utility (benefit), and the rejectability membership function is an in-utility (cost). For driving,

[16] The SB agent states are sub-states of $\theta_{\text{RB}} \supset \theta_{\text{SB}}$. For u_{RB} =BA the SB state is $\theta_{\text{SB}} = [T_c^{-1}, T_h]^T$, for u_{RB} =TR the SB state is $\theta_{\text{SB}} = [T_c^{-1}, T_h]^T$, and for u_{RB} =SR the SB state is $\theta_{\text{SB}} = v_A$.

global information is necessary to determine if a chosen speed not only moves you toward your destination expediently (a benefit) but also without incident (a cost). However, from experience drivers learn to recognize that some conditions are locally expedient but may not be globally safe (e.g., traveling fast may be expedient but may also cause an accident), and that some conditions are locally safe but may not be globally expedient (e.g., parking your car may prevent a collision but may also prevent you from reaching your destination). Thus, it can be argued that RB agents possess task specific values based upon local information[17] (such as speed and headway) that represent global (KB agent) consequences (such as safety and expediency). Such values are represented by the accuracy and rejectability (utility and in-utility) membership functions. Using these observations, we associate the global goals of expediency and safety with local values based on T_c^{-1} and T_h, respectively.

Consider the decision to switch to BA from either TR or SR. For such a switch, the sub-state $[T_c^{-1}, T_h]^T$ of θ_{RB} can be used to determine when nominal SB behavior is satisficing (i.e., when the sub-state $[T_c^{-1}, T_h]^T$ is such that $\theta_{\mathrm{RB}} \notin S_b(u_{\mathrm{RB}}), u_{\mathrm{RB}} \in \{\mathrm{SR}, \mathrm{TR}\}$ then neither TR nor SR should be used). A small T_c^{-1} (small relative velocity) indicates that vehicle A is appropriately following vehicle B such that the driver is traveling at an expedient speed (driving as fast as possible without risking incident). A small T_h indicates that the relative distance R between vehicles is small given v_A, which is associated with danger even if expedient (any change in the preceding vehicle speed or any error in the perceptual state estimate can produce a dangerously low time to collision). Thus, low (including highly negative) T_c^{-1} has high benefit, and low (but positive, negative headway implies a crash) T_h has high cost.

By independently describing and comparing values based on time to collision and time headway, we make explicit the tradeoff between expediency and safety. Such an approach seems consistent with intuition regarding driving: keep time to collision large so that you travel as fast as possible and still leave yourself time to get out of a dangerous situation; and avoid small time headways so that you avoid encountering dangerous situations and operate within acceptable limits posed by reaction time and uncertainty.

4.4 Empirical Estimates

We now describe how (partial) estimates of $\mu_A(u_{\mathrm{RB}}, \theta_{\mathrm{RB}})$ and $\mu_R(u_{\mathrm{RB}}, \theta_{\mathrm{RB}})$ can be identified[18] can be obtained from average (across all subjects) empirical data. In this identification, our objective is to find substates that trigger active braking. We therefore distinguish between nominal behavior $u_{\mathrm{RB}} \in \{\mathrm{SR}, \mathrm{TR}\}$ and active braking behavior $u_{\mathrm{RB}} = \mathrm{BA}$. Our goal is thus to find when $\theta_{\mathrm{RB}} \notin S_b(u_{\mathrm{RB}})$ for $u_{\mathrm{RB}} \in \{\mathrm{SR}, \mathrm{TR}\}$. Nominal operating conditions occur when the brake pedal is not pressed. For both nominal and braking conditions, we select representative sample points from each experimental trial and create two sets of $[T_c^{-1}, T_h]^T$ points: one set for nominal conditions, denoted NOM, and one set for braking conditions, denoted BRK. For trials when subjects actively brake, the sub-state(s)[19] $[T_c^{-1}, T_h]^T$ when braking is initiated is included in BRK, and the sub-state(s) $[T_c^{-1}, T_h]^T$ when braking is terminated is included in NOM; for trials when subjects do not brake, the initial sub-state $[T_c^{-1}, T_h]^T$ in the trial is included in NOM; and for trials where subjects only brake (by anticipating the cut-in and then coming to a stop), the initial sub-state $[T_c^{-1}, T_h]^T$ in the trial is included in BRK.

For notational purposes in the subsequent sections, let $N(T = \tau|\mathsf{CONDITION})$ denote the cardinality of the set of points $T = \tau$ given CONDITION. For example, $N(T_c^{-1} = \tau|\mathsf{NOM})$ is the number of points in the set $\{\theta_{\mathrm{RB}} \in \mathsf{NOM} : T_c^{-1} = \tau\}$. We define $\mu_A(\theta_{\mathrm{RB}})$ and $\mu_R(\theta_{\mathrm{SB}})$ as functions of θ_{RB} but not u_{RB}, and use these membership functions to determine the boundary between braking and nominal behaviors. These "partial" membership functions represent the value of the perceptual states T_c^{-1} and T_h in relation to the goals of expedient but safe driving.

Accuracy Under nominal conditions ($\theta_{\mathrm{RB}} \in \mathsf{NOM}$), relative velocity must be considered acceptable to the driver whence the distribution of T_c^{-1} under nominal conditions is an observable entity that provides information about what is accurate. Clearly, if $T_c^{-1} = \tau_2$ is accurate, then $\tau_1 < \tau_2$ must be at least as

[17]Goals and values exist in different temporal worlds; goals are global and values are local instantiations of goals triggered by perceptual cues. For example, for car-following the global goals are to reach a destination safely and expediently, and the local values are determined by current and future perceptual states.

[18]Our approach is slightly oversimplified because braking and acceleration characteristics are confounded by perceptual thresholds.

[19]In several trials, subjects initiate braking, establish large T_h, and re-initiate the CC system. When vehicle B is traveling slower that the set CC speed, this causes vehicle A to close on vehicleB and thereby causes the driver to re-apply braking.

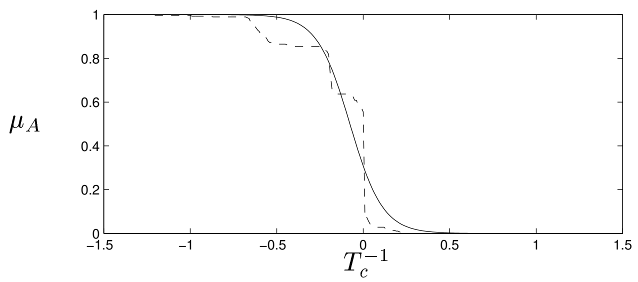

Figure 19. Actual (dashed line) and approximated (solid line) accuracy membership functions.

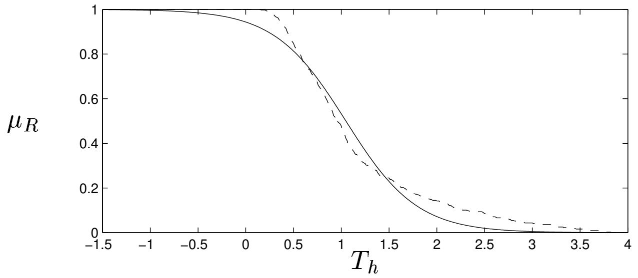

Figure 20. Actual (dashed line) and approximated (solid line) rejectability membership functions.

accurate. This monotonicity property facilitates the computation of the accuracy function as the cumulative distribution function

$$\mu_A(T_c^{-1} = \tau) = 1 - F_{T_c^{-1}}(\tau|\mathsf{NOM}) = 1 - \frac{N(T_c^{-1} \le \tau|\mathsf{NOM})}{N(T_c^{-1} \le \infty|\mathsf{NOM})}.$$

For classification purposes, we fit (via least squares) a sigma function of the form $1/e^{(-a\tau+b)}$ to $\mu_A(\cdot)$ yielding the membership function shown in Figure 19.

Rejectability When braking is initiated ($\theta_{\mathrm{RB}} \in \mathsf{BRK}$), time headway values must be considered unacceptable whence the distribution of time headways when the driver initiates braking is an observable entity that provides information about what is rejectable. Clearly, if $T_h = \tau_2$ is rejectable then $\tau_1 < \tau_2$ must be at least as rejectable. This monotonicity property facilitates the computation of the rejectability function as the cumulative distribution function

$$\mu_R(T_h = \tau) = 1 - F_{T_h}(\tau|\mathsf{BRK}) = 1 - \frac{N(T_h \le \tau|\mathsf{BRK})}{N(T_h \le \infty|\mathsf{BRK})}.$$

For classification purposes, we fit (via least squares) a sigma function of the form $1/e^{(-a\tau+b)}$ to $\mu_R(\cdot)$ yielding the membership function shown in Figure 20.

4.5 Classification Results

For the RB agent to determine switching from one SB agent to another, it is necessary to identify when $u_{\mathrm{RB}} \notin S_b(\theta_{\mathrm{RB}})$. Using $\mu_A(T_c^{-1})$ and $\mu_R(T_h)$ we can construct the set of states $\mathcal{S}_b = \{\theta_{\mathrm{RB}} : \mu_A(T_c^{-1}) \ge b\mu_R(T_h)\}$ that support nominal behavior, and the set of states $\mathcal{S}_b^c = \{\theta_{\mathrm{RB}} : \mu_A(T_c^{-1}) < b\mu_R(T_h)\}$ (superscript c denotes complement) that do not support nominal behavior. If $u_{\mathrm{RB}} \in \{\mathrm{TR}, \mathrm{SR}\}$ and $\theta_{\mathrm{RB}} \in \mathcal{S}_b^c$ then $\theta_{\mathrm{RB}} \notin S_b(u_{\mathrm{RB}})$. Thus, the line $\mu_A(T_c^{-1}) = b\mu_R(T_h)$ determines when behavior must be switched from nominal to braking[20].

Given the empirically derived membership functions, we can determine the boundary between nominal and braking behaviors as a function of b by finding the perceptual states θ for which $\mu_A(T_c^{-1}) = b\mu_R(T_h)$. This is illustrated in Figure 21 for the data gathered in the simulator experiment, where $\circ$ indicates $\theta_{\mathrm{RB}} \in \mathsf{NOM}$ and $\times$ indicates $\theta_{\mathrm{RB}} \in \mathsf{BRK}$. To the northwest of the line, BA is satisficing but TR and SR are not, and to the southeast of the line TR and SR (and, perhaps, BA) are satisficing. Classification can be performed by finding the value of b that optimally separates braking from nominal behavior. To determine optimality, let $N[H]$ denote the cardinality of set H, and consider the following three performance indices

$$\begin{aligned} J_1(b) &= \frac{N\left[(\mathsf{NOM} \cap \mathcal{S}_b^c) \bigcup (\mathsf{BRK} \cap \mathcal{S}_b)\right]}{N\left[\mathsf{NOM} \bigcup \mathsf{BRK}\right]} \\ J_2(b) &= \frac{N\left[\mathsf{NOM} \cap \mathcal{S}_b^c\right]}{N\left[\mathsf{NOM}\right]} \end{aligned}$$

[20]This line does not indicate when braking behavior should be switched to car following (TR) behavior. In general, there are many perceptual states for which SR/TR and BA behaviors are simultaneously satisficing. As an example, drivers may initiate the brakes to avoid a collision and continue to press even when $T_c^{-1} < 0$ and T_h is close to T_h^*. If, by contrast, a vehicle cuts-in with $T_c^{-1} < 0$ and T_h is close to T_h^* the driver may not brake but instead use engine braking or simply wait until the cut-in vehicle establishes a safe lead distance. Thus, this state $\boldsymbol{\chi}$ is satisficing for both BA and SR/TR

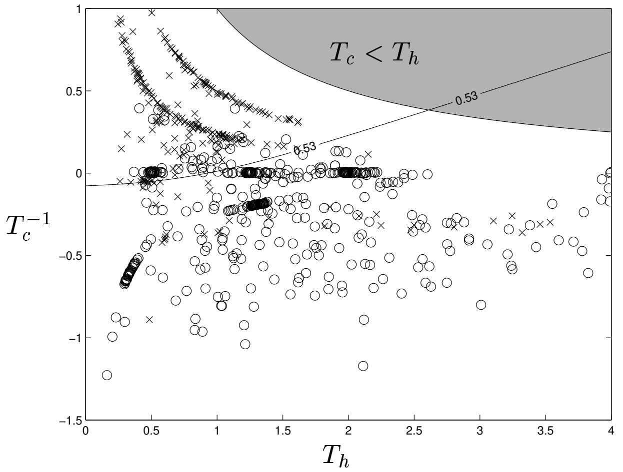

Figure 21. Scatter plot of nominal and braking perceptual states.

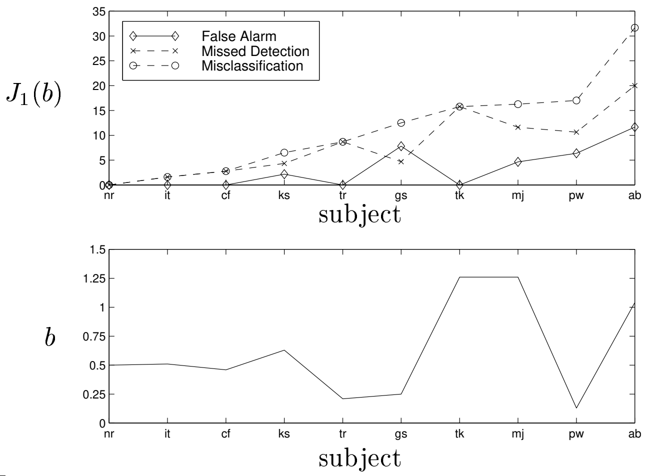

Figure 22. Classification results per subject.

$$J_3(b) \quad = \quad \frac{N\left[\mathsf{BRK} \cap \mathcal{S}_b\right]}{N\left[\mathsf{BRK}|\right]}.$$

Intuitively, $J_1(b)$ is the percentage of trials that are incorrectly classified (i.e., the total number of $\circ$'s above the line plus the total number of $\times$'s below the line) , $J_2(b)$ is the percentage of nominal trials that are incorrectly classified as braking (i.e., number of $\circ$'s above the line), and $J_3(b)$ is the percentage of braking trials that are incorrectly classified as nominal (i.e., number of $\times$'s below the line). The value $b = 0.53$ is the minimax value $b = \arg\min_{b \geq 0} \max\{J_1(b), J_2(b), J_3(b)\}$ which attempts to balance the percentage of misclassifications ($J_1(b)$), false alarms ($J_2(b)$), and missed detections ($J_3(b)$). The value $b = 0.20$ minimizes the number of samples misclassified $b = \arg\min_{b\geq 0} J_1(b)$. The classification results for the different values of b are shown in Table 6 and indicate that, on the average, over 85% of samples are correctly classified.

b	% misclassified	% false braking	% missed braking
0.20	10.04	1.95	8.09
0.53	13.25	8.37	4.88

Table 6. Classification accuracies for different values of b.

Given the average membership functions $\mu_A(T_c^{-1})$ and $\mu_R(T_h)$, we can find the optimal (with respect to $J_1(b)$) b value for each subject. This value of b and the corresponding classification percentage is shown in Figure 22[21].

Note that the line separating braking from nominal behavior in Figure 21 is curved. Thus, although we speak loosely and say that $T_c^{-1} \geq 0$ (range between vehicles decreasing) triggers braking, it is more precise to say that braking is triggered by a combination of time to collision and time headway measurements. This is consistent with observations of other researchers such as [20].

5 Experiment 3: Modeling Braking Responses

The key points in this section are as follows:

- Drivers display significantly different steady state time headways[22].

[21] Note that subject `ab` was omitted from computing the average partial membership functions because this subject exhibited very unusual behavior. This subject drove race cars in Italy and, presumably, drove with one foot on the brake and one foot on the accelerator. Consequently, the thresholds we used to distinguish between braking and nominal behavior were ambiguous and unusual classification resulted.

[22] An interesting area of future research is to determine if behaviors of drivers with low steady state time headways correlate with high b parameters in SDT.

- SDT can be used to distinguish between nominal and braking behaviors in real vehicles.
- Observed SB behaviors, which are characterized by smooth counter clockwise trajectories in perceptual space, can be closely matched by the parameter rich MPC formalism.
- Some MPC parameters exhibit a sensitivity to initial T_c^{-1} and T_h values.

5.1 Experiment Description

In the experiment, two vehicles drive in the same lane on a closed test track. The subject drives vehicle A which follows vehicle B. The drivers in vehicles A and B are required to maintain an assigned speed $v_A(0)$ and v_B until a chime rings in vehicle A's car. When the chime rings, the driver of vehicle A is to establish a natural following distance (i.e., drive as if vehicle B had just cut-in to vehicle A's lane) while vehicle B maintains a constant speed. Two professional drivers participated as subjects (driver's of vehicle A), and two values of speed were selected for vehicle B. The selected relative distances $R(0)$ that trigger the chimes and the other experimental conditions are tabulated in Figure 7. Measurements, taken at a sampling rate

$v_B = 50$ km/hr		$v_B = 70$ km/hr	
$v_A(0)$ (km/hr)	$R(0)$ (m)	$v_A(0)$ (km/hr)	$R(0)$ (m)
50	5	70	25
50	15	70	35
50	25	70	40
60	15	80	25
60	25	80	35
60	30	80	40
70	25	90	25
70	35	90	35
70	40	90	40
80	25	100	25
80	35	100	35
80	40	100	40

Table 7. Experimental trials for braking dynamics. For $v_B = 70$ and $v_A(0) = 70$, no active braking occurs for either subject, and for $v_B = 70$, $v_A(0) = 80$, and $R(0) = 40$ subject 2 does not actively brake. For subject 1, errors in the measurement methods invalidated observations with $v_B = 50$ and both $v_A(0) = 70$, $R(0) = 25$ as well as $v_A(0) = 80$, $R(0) = 25$.

of 50 ms, include R, v_A, brake pressure β, and throttle opening angle α. To compute T_c^{-1}, it is necessary to determine v_R. We do this by taking the temporal difference of R using MATLAB's `gradient` function, negating (recall the sign convention $v_R = -\dot{R}$ from the introduction), dividing by the sample time T_s, and filtering to remove noise. In practice, we use a fourth order discrete Butterworth low-pass filter with cut-off frequency of $0.05/T_s$.

5.2 Results: Brake Initiation via SDT

From the experimental data, three observations are worth noting. First (see Figures 23-24) subjects establish natural following distances by generating smooth counterclockwise trajectories in perceptual space. This supports the observations made in the first three experiments. Second (again, see Figures 23-24), subject 1 and subject 2 establish steady-state (i.e., $v_R \approx 0$) behavior at different values of T_h^*. Not only is this true for the initial conditions shown in the figures, but also for every other initial condition. In fact, there is a significant ($P \approx 2 \times 10^{-7}$) T_h^* difference between drivers. For subject A the mean terminal headway is $T_h^* = 1.47$, and for subject B the mean terminal headway is $T_h^* = 2.01$. Interestingly, there are no significant within subject T_h^* differences for different $v_A(0)$ or v_B conditions. Thus, we find evidence that T_h influences braking dynamics independently of v_A (see discussion in Section 3.3). Third (see Figures 25-26), we can classify nominal and braking behaviors using SDT. To perform this classification, we construct $\mu_A(T_c^{-1})$ and $\mu_R(T_h)$ (using all nominal conditions and all braking conditions, respectively, rather than just initial nominal/braking conditions), and then find b that minimizes the misclassification. The results indicate one false alarm ($\circ$ above the line) and no missed detections ($\times$ below the line).

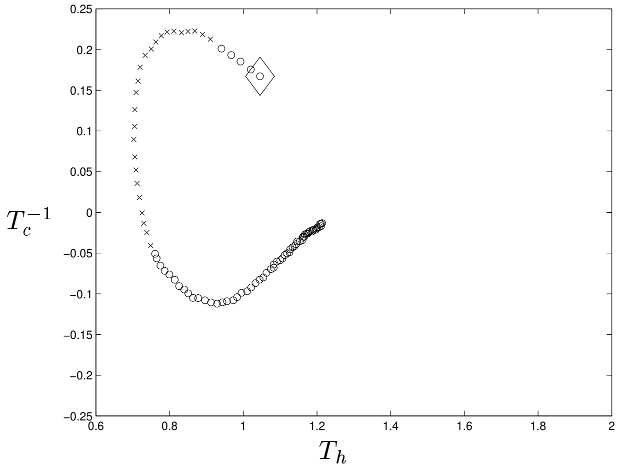

Figure 23. Perceptual phase plane trajectories for subject 1 with initial conditions $(T_h(0), v_R(0)) \approx (1\text{s}, 20\text{km/hr})$.

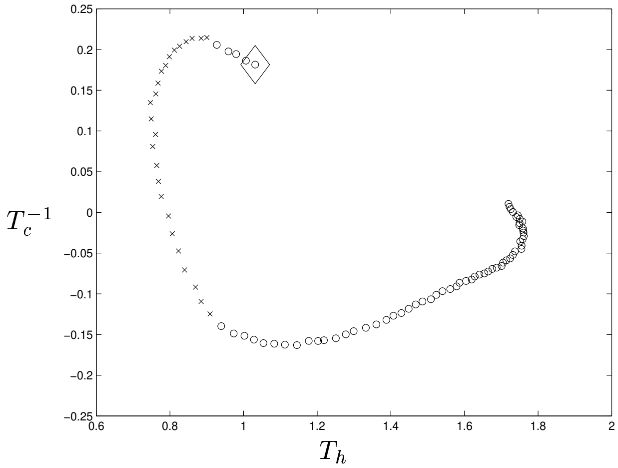

Figure 24. Perceptual phase plane trajectories for subject 2 with initial conditions $(T_h(0), v_R(0)) \approx (1\text{s}, 20\text{km/hr})$.

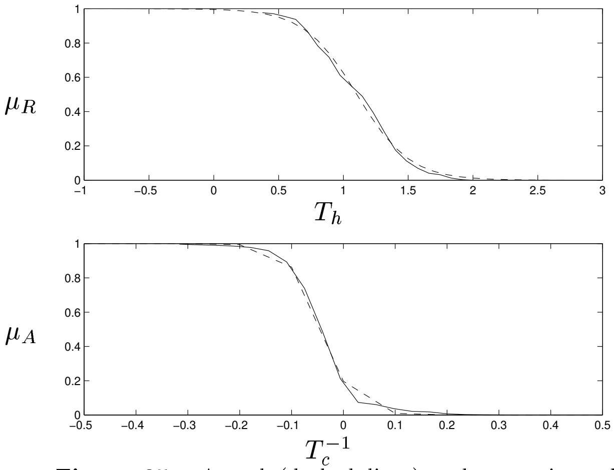

Figure 25. Actual (dashed lines) and approximated (solid lines) membership functions.

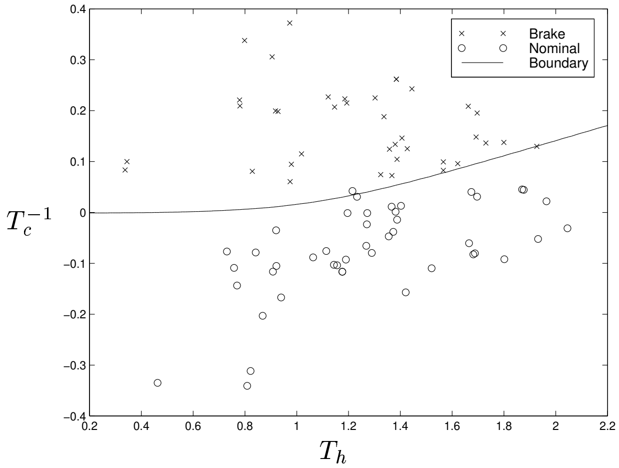

Figure 26. Perceptual classification.

5.3 Results: Brake Dynamics via MPC

5.3.1 Vehicle Model

To apply model predictive control, we require a discrete time dynamical model of vehicle dynamics. For our purposes, it is sufficient to determine how brake pressure β ($\beta > 0$ implies braking) and throttle opening angle α ($\alpha > 0$ implies throttle open) affect velocity v_A whence we consider the simple linear model

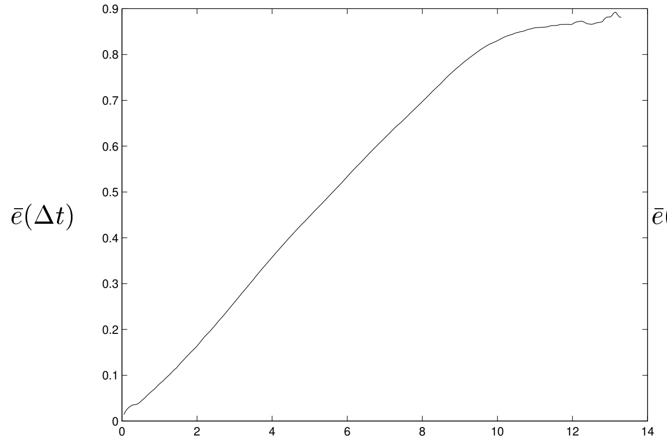

Figure 27. Accumulation of linear model prediction errors for subject 1 under $v_B = 70$, $v_A(0) = 90$, and $R(0) = 25$.

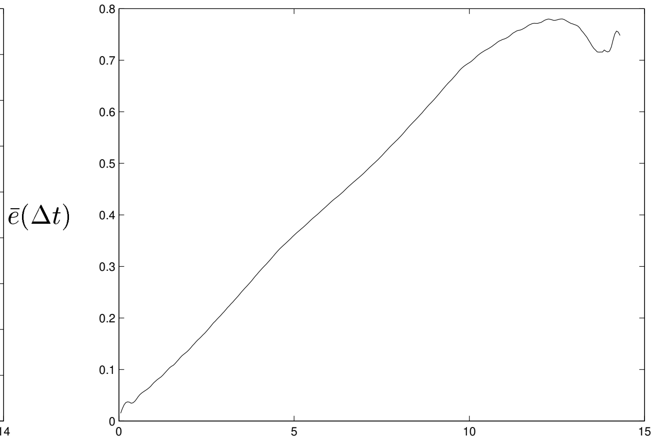

Figure 28. Accumulation of linear model prediction errors for subject 2 under $v_B = 70$, $v_A(0) = 90$, and $R(0) = 25$.

$$v_A(k+1) = Av_A(k) + B \begin{bmatrix} \beta(k) \\ \alpha(k) \end{bmatrix}. \tag{12}$$

Fitting a curve to the measured velocities for each trial, eliminating outliers, and averaging parameter values yields (note that there is no significant difference of parameter values as a function of velocity or driver; this is expected unless one or both drivers excite nonlinear portion of "real" system.)

$$A = 0.9993 \quad B = [-0.0372, 0.0035].$$

The purpose of identifying (12) is to generate predictions of future vehicle behavior as a function of control $u_A = [\beta, \alpha]^T$. For any measured velocity $v_A(t)$ and control sequence $u_A(t), \ldots, u_A(t+\Delta t-1)$ we can estimate the resulting sequence of velocities $\hat{v}_A(t+1), \ldots, \hat{v}_A(t+\Delta t)$ from (12). Given this estimate, we can quantify the error between the measured velocity v_A and estimated velocity $\hat{v}_A$ at time $t + \Delta t$

$$e(t, \Delta t) = |v_A(t+\Delta t) - \hat{v}_A(t+\Delta t)|.$$

For the observed velocity sequence $v_A(t_0), \ldots, v_A(t_f)$, we can compute the average error of the linear model as a function of prediction time Δt

$$\bar{e}(\Delta t) = \frac{\sum_{t=t_0}^{t_f-\Delta t} e(t, \Delta t)}{t_f - \Delta t}.$$

This average error function is plotted in Figures 27-28 and shows that, for prediction horizons on the order of one to two seconds, the cumulative model prediction error is less than[23] 1% of v_A. Given these acceptable bounds, we can define the state $\boldsymbol{x} = [d_A, v_A, d_B, v_B]$, with d_A and d_B the distances (in relative coordinates) of vehicles A and B, respectively. The discrete time dynamics can then be identified as

$$\begin{aligned} \boldsymbol{x}(k+1) &= \boldsymbol{g}(\boldsymbol{x}_k, u_A, u_B) \\ &= \begin{bmatrix} 1 & T_s & 0 & 0 \\ 0 & A & 0 & 0 \\ 0 & 0 & 1 & T_s \\ 0 & 0 & 0 & A \end{bmatrix} \boldsymbol{x}(k) + \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ B & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \\ \mathbf{0} & B \end{bmatrix} \begin{bmatrix} u_A \\ u_B \end{bmatrix}, \end{aligned} \tag{13}$$

[23] Actually, the linear model is a much better fit for braking than acceleration.

where $T_s = 0.05$ (s) is the sample time of the measured data and where we have assumed that vehicle B has the same dynamics as vehicle A. Given (13), we can identify the perceptual state dynamics $\boldsymbol{\chi}(k+1) = \hat{\boldsymbol{f}}(\boldsymbol{\chi}(k), u_A, u_B)$ using the definitions of T_c^{-1} and T_h in (1) with range $R = d_B - d_A$. To simplify the complexity of the MPC controller, we use the sample time[24] of $T_s^p = 0.5$ (s) (superscript p for predictive) and require that u_A be constant over the prediction interval $t \in [0, T_s^p]$.

5.3.2 Parameter Search

Since we are interested in responding to cut-in events, we are interested in both TR and BA behaviors whence we consider the perceptual state $\theta_{\text{SB}} = \boldsymbol{\chi} = [T_c^{-1}, T_h]^T$. The model predictive control $u_A = \boldsymbol{\pi}(\boldsymbol{f}, \boldsymbol{\chi}_0)$ is obtained by minimizing the cost function (7) with terminal and incremental costs given by (9)-(10). We can identify the parameter set P, Q, S, and S_Δ by generating braking behavior via MPC and comparing this to observed braking behavior. In this identification, our first step is to select the planning horizon. For our purposes, we employ a nominal two second planning horizon whence $N = 2/T_s^p = 4$. The second step is to identify the desired perceptual state. We use the characteristic goals of driving $T_c^{-1} \to 0$ and $T_h \to T_h^*$, where T_h^* is selected as the terminal headway for each experimental condition.

The third step is to select the appropriate perceptual state. We suppose that time to collision and time headway independently contribute to the braking profile whence we restrict attention to diagonal $P = Q = \begin{bmatrix} q_{T_c^{-1}} & 0 \\ 0 & q_{T_h} \end{bmatrix}$. Let $\Phi = [q_{T_c^{-1}}, q_{T_h}, S, S_\Delta, u_B]$, and let $\hat{u}_A = \pi(\hat{\boldsymbol{f}}, \boldsymbol{\chi}_0, \Phi)$ denote the MPC law parameterized by Φ using the perceptual dynamics model $\hat{\boldsymbol{f}}$ derived in the previous section. We can identify Φ by defining a measure of fitness $\mathcal{J}$ for $\pi(\hat{\boldsymbol{f}}, \boldsymbol{\chi}_0, \Phi)$ with respect to observed driver behavior. To do so for any initial state $\boldsymbol{\chi}_0 = \hat{\boldsymbol{\chi}}(t+\tau-1)$, we consider the sequence of states $\hat{\boldsymbol{\chi}}(t+\tau)$ that result from applying control law $\pi(\hat{\boldsymbol{f}}, \boldsymbol{\chi}_0, \Phi)$. We compare this sequence to the sequence of measured states $\boldsymbol{\chi}(t_0), \ldots, \boldsymbol{\chi}(t_f)$, and define $\mathcal{J}$ as

$$\mathcal{J}(\Phi) = \sum_{t=t_0}^{t_f} \left(\boldsymbol{\chi}(t) - \hat{\boldsymbol{\chi}}(t)\right)^2.$$

We select t_0 as the point in time when braking is initiated, and t_f as the point in time when braking is terminated[25]. Minimizing $\mathcal{J}(\Phi)$ with respect to Φ then gives the optimal parameter estimates

$$\Phi^* = \arg\min \mathcal{J}(\Phi).$$

To perform this search, we use MATLAB's `fmins` simplex search function with initial conditions $\Phi_0 = [12, 4, 4, 4, 5]^T$.

5.3.3 Parameter Results

In Figures 29-32 we plot the perceptual phase plane trajectories and brake pressure time histories for both the observed and MPC-generated behaviors given two experimental conditions. Note the similar behavior exhibited in both perceptual state space trajectories and in brake pressure histories. Given the set of Φ^* estimates for each set of experimental conditions, it is necessary to quantify the usefulness of these parameters.

To do so, we can analyze Φ's sensitivity on experiment conditions. Our first step in performing this analysis is to examine the ratio of standard deviation to mean for each element in Φ. This ratio $\boldsymbol{r} = [r_{Q_{T_c^{-1}}}, r_{Q_{T_h}}, r_S, r_{S_\Delta}, r_{u_B}]^T = [0.16, 0.16, 0.69, 0.29, 0.21]^T$. From this sensitivity analysis, we are motivated to look for dependencies of S on experimental conditions. This sensitivity is graphically depicted in Figure 33. A similar analysis was performed for for a best fit linear model with state $\boldsymbol{x} = [v_A, a_A, v_B, a_B]^T$, with a_A representing acceleration, and dynamics $\boldsymbol{x}(t+1) = A\boldsymbol{x}(t) + B[u_A(t)u_B(t)]$. The results are similar qualitative and quantitative (in terms of sum of squared trajectory errors) fits to the observed perceptual phase trajectories. Moreover, the best fit parameters display a similar sensitivity to initial T_c^{-1} and T_h values.

[24] Note that we are using two different sample times: T_s^p is the sample time used to determine the MPC law $u = \boldsymbol{\pi}(\boldsymbol{f}, \boldsymbol{\chi})$, and T_s is the rate at which u is updated.

[25] In keeping with the multi-agent description, brake termination triggers the TR SB controller. This controller may have a different parameter set Φ.

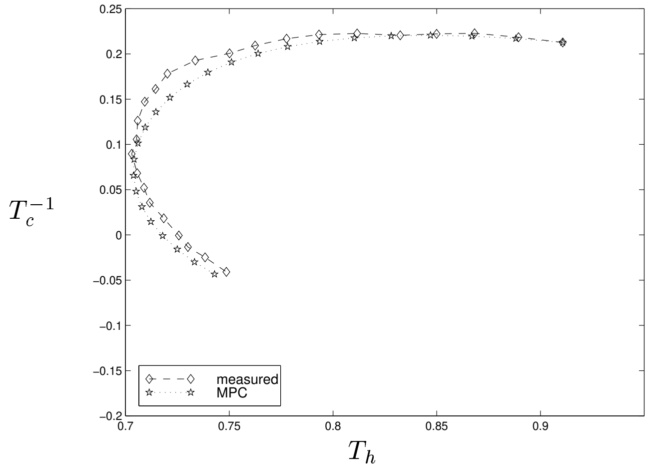

Figure 29. Perceptual trajectories for subject 1 under $v_B = 70$, $v_A(0) = 90$, and $R(0) = 25$.

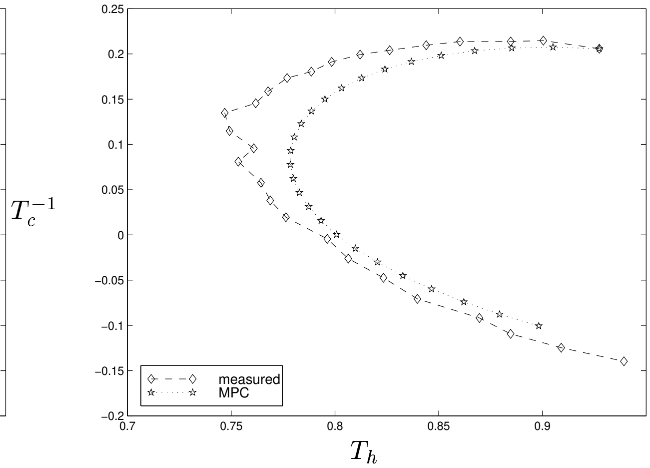

Figure 30. Perceptual trajectories for subject 2 under $v_B = 70$, $v_A(0) = 90$, and $R(0) = 25$.

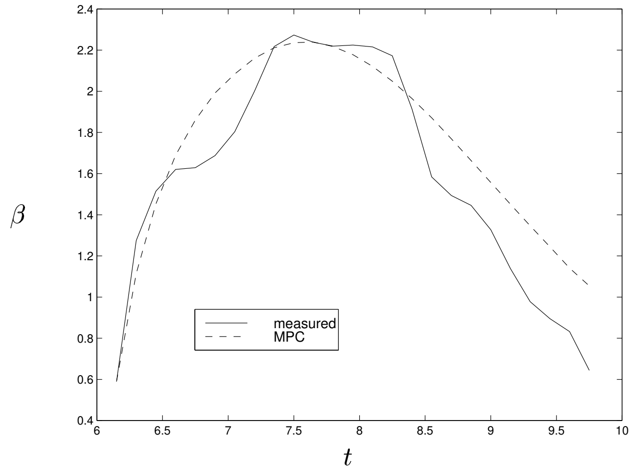

Figure 31. Brake pressure history for subject 1 under $v_B = 70$, $v_A(0) = 90$, and $R(0) = 25$.

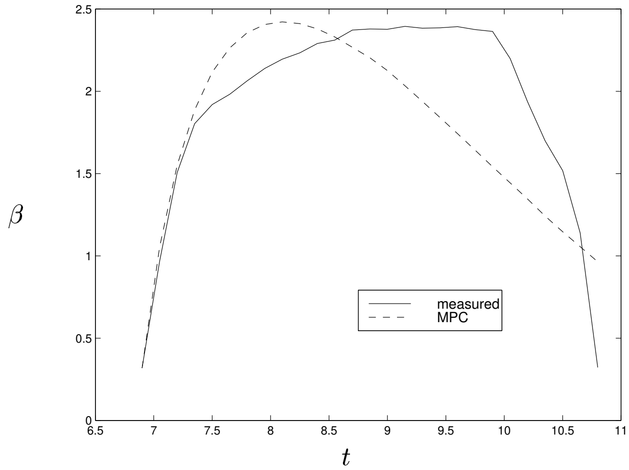

Figure 32. Brake pressure history for subject 2 under $v_B = 70$, $v_A(0) = 90$, and $R(0) = 25$.

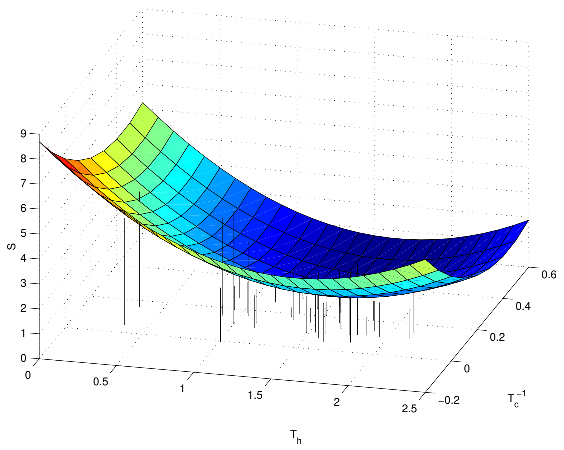

Figure 33. S sensitivity to initial $= \boldsymbol{\chi} = [T_c^{-1}, T_h]^T$, and best fit to quadratic function of $\boldsymbol{\chi}$.

5.3.4 Tradeoffs in MPC-based Longitudinal Control

For MPC, the ratio of Q to S represents the tradeoff between the goal of reaching $\boldsymbol{\chi}^*$ and the goal to maintain comfort. Let $\rho(P)$ be the maximum eigenvalue of matrix P. Then, the ratio $\ell(Q)/\max(\rho(S), \rho(S_\Delta))$ indicates the "comfort level" of the MPC control. Figure 34 illustrates how the braking profile changes as a function of ℓ. For high ℓ, braking pressure is high and the perceptual trajectory rapidly approaches $T_c^{-1} = 0$ and $T_h \to T_h^*$; even though v_B is nearly constant the perceptual trajectory is nearly clockwise. Such values of ℓ are likely to be perceived as safe, but uncomfortable. For low ℓ, braking pressure is lower and convergence to $T_c^{-1} = 0$ and $T_h \to T_h^*$ is much slower. Such values of ℓ are likely to be perceived as comfortable, but unsafe. The effect of $\boldsymbol{u}$ is an area of continuing research.

Given that Φ^* is sensitive to initial $\boldsymbol{\chi}(0)$ conditions, it is apparent that ℓ also depends on initial perceptual conditions. Although we are continuing to analyze the MPC approach, it appears that the MPC parameters must be scheduled according to $\boldsymbol{\chi}(0)$. Because of this sensitivity, it is desirable to consider other approaches that can more effectively encode state-dependent parameters.

6 Future Work, Summary, and Conclusions

6.1 Future Work

Based on the theoretical framework of Sections 1-2 and the experimental evidence in Sections 3-5, we can identify several areas of future research. Beginning at the bottom of the KRS hierarchy, appropriate models of SB control will continue to be developed. Of immediate interest is the construction of models without the parametric sensitivity to initial conditions that was exhibited by the MPC approach. To perform such modeling, we will explore the use of both MPC with superior vehicle dynamics models as well as neuro-fuzzy control.

In the middle of the KRS hierarchy, it is desirable to explore the coordination of multiple RB agents. It seems likely that some of this coordination will be governed by KB agents and that other coordination will be distributed among the RB agents. For distributed coordination, it is desirable to consider methods to fuse the behaviors produced by the RB task managers to generate sophisticated behaviors. Robust SDT [17] and generalized potential field [9, 22, 37] methods are both candidates for such behavior fusion. For coordination directed by KB agents, methods both for scheduling attention as well as for modeling communication channels must be formulated.

At the top of the KRS hierarchy, we note that much of driving does not require KB intervention. The tem-

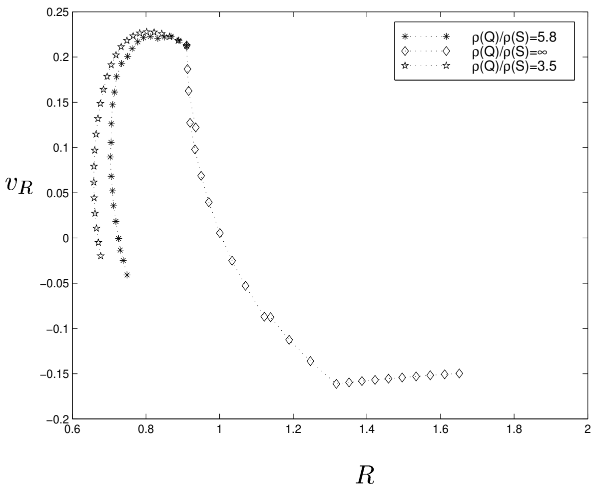

Figure 34. Braking profile sensitivity to ℓ.

poral dynamics of much of driving prevent the usage of slow conscious deliberation. Perhaps the exceptions to this are in navigation and learning tasks.

Finally, it is important that the impact of learning on the multiple mental model society be investigated. Such learning includes SB control improvement, RB task acquisition, and KB goal definition and adaptation.

As we continue to develop multiple agent models of driver behavior, our next step will be to study how lateral and longitudinal agents interact to generate behavior. Of particular interest is a model of how drivers negotiate curves in the presence of other traffic, and how drivers change lanes in response to other vehicle behaviors.

6.2 Summary

In this paper, we have outlined a theoretical model of driver decision-making and behavior generation that utilizes a three level hierarchy of knowledge-based, rule-based, and skill-based agents. We use SDT to model and describe RB task management for longitudinal control in response to a vehicle cut-in, and MPC to emulate SB braking response to a vehicle cut-in.

Experiments have been conducted that support the following hypotheses about driving behavior:

- The perceptual sub-state $\chi = [T_c^{-1}, T_h]^T$ can be used to efficiently classify behavior into braking and nominal categories. Braking is initiated (ignoring reaction) when $T_c^{-1} > 0$ ($v_R > 0$). Conversely, a driver is likely to accelerate when $T_c^{-1} < 0$ ($v_R < 0$). Thus, dividing driver behavior into active braking and nominal (not-active) braking produces a division roughly at $T_C^{-1} = 0$ ($v_R = 0$). Additionally, the precise T_c^{-1} threshold increases as T_h increases; thus, a driver's decision to brake is a function of both safety and expediency.

- Classification of active braking and nominal behavior (i.e., speed or time headway regulation) in both simulator and physical driving scenarios can be successfully performed using SDT with physically interpretable set membership functions.

- Drivers appear to first establish infinite time to collision ($T_c^{-1} \leq 0$) and then drive the system to a desired time headway T_h^*. This corresponds to controlling first to establish safety and second to reestablish an expedient safety margin. Acceptable behaviors are characterized by a counterclockwise perceptual phase trajectory, and unacceptable trajectories by a clockwise trajectory. Phase plane trajectories in perceptual space must not only be counterclockwise, they must also be smooth.

- Observed SB behaviors, which are characterized by smooth counterclockwise trajectories in perceptual space, can be closely matched by the parameter rich, MPC formalism. Some MPC parameters exhibit a sensitivity to initial T_c^{-1} and T_h values.
- Drivers display significantly different desired (steady state) time headways.

6.3 Conclusions

A driver interprets and responds to sensory input using mental models. The mental model selects appropriate perceptual cues and interprets these cues based on its intended goal. Following the example of multi-agent intelligent systems, multiple mental models can be organized into a multi-resolutional society with knowledge-based, rule-based, and skill-based controllers. Rule-based task managers select skill-based controllers depending on whether the perceptual state is satisficing for the controller. Such task management can be emulated by SDT with the perceptual states T_c^{-1} and T_h used to detect satisficing performance. Skill-based task controllers can be emulated by perceptually scheduled MPC that accomplishes the driver goals of driving $T_c^{-1} \to 0$ and $T_h \to T_h^*$.

A Attentional Updating

To effectively coordinate mental models, communication within the society is necessary. Child agents communicate their current state and their attentional requirements to their parents (bottom up communication), and parent agents allocate this attentional resource and dictate switching between child agents (top down communication). Such communication is represented in Figures 2-4 by directional arrows. Tasks associated with high workload and high perceptual bandwidth demand high attentional resources, and tasks associated with low workload require low attentional resources. It is necessary for SB controllers and RB task managers to communicate (from the bottom-up) such requirements to their parents [55].

Beginning at the bottom with SB agents, there exists a dynamic relation between past $\theta_{\mathrm{SB}}(k-1)$ and current $\theta_{\mathrm{SB}}(k)$ as a function of SB action u_{SB}

$$\theta_{\mathrm{SB}}(k) = f(\theta_{\mathrm{SB}}(k-1), u_{\mathrm{SB}}(k-1), \zeta(k-1)), \tag{14}$$

where $\zeta(k-1)$ represents a disturbance such as another driver's action. When the SB agent is engaged (attention is held), current estimates of the perceptual state are obtained from

$$\hat{\theta}_{\mathrm{SB}}(k) = \theta_{\mathrm{SB}}(k) + \eta(k), \tag{15}$$

where $\eta(k)$ represents sensory-perception noise, and $\theta_{\mathrm{SB}}(k)$ represents the "true" perceptual state. When an SB agent is disengaged (attention is not held), estimates of the current perceptual state are obtained through open loop predictions (i.e., no sensory perception) obtained from an internal model of (14)

$$\hat{\theta}_{\mathrm{SB}}(k) = \hat{f}(\hat{\theta}_{\mathrm{SB}}(k-1), u_{\mathrm{SB}}(k-1), \hat{\zeta}(k-1)), \tag{16}$$

Continuing from the bottom toward the top, an RB agent amalgamates relevant $\hat{\theta}_{\mathrm{SB}}$ to form $\hat{\theta}_{\mathrm{RB}}$, and then propagates the error covariance $P(k)$ of the estimation error $\tilde{\theta}_{\mathrm{RB}}(k)$

$$\tilde{\theta}_{\mathrm{RB}}(k) = \theta_{\mathrm{RB}}(k) - \hat{\theta}_{\mathrm{RB}}(k) \qquad P(k) = E\tilde{\theta}_{\mathrm{RB}}(k)\tilde{\theta}_{\mathrm{RB}}(k)^T.$$

Without current perceptual measurements (i.e., with open loop estimates (16)), the covariance matrix grows until eventually the boundary of this matrix overlaps a perceptual region that is not satisficing to the current RB agent. By communicating the rate at which $\theta_{\mathrm{SB}}(k)$ changes (and hence how $\theta_{\mathrm{RB}}(k)$ changes and $P(k)$ grows), the SB agent communicates its need for attentional resources required to accomplish its assigned task. Given this rate information, the RB agent determines the amount of time available before the range of possible errors is unacceptable, and communicates this time to the KB agent.

At the top, a KB agent requires an estimate of when attention might be needed again prior to switching attention from one RB level task to another. This amount of time is communicated from the RB agents to the KB agent who then schedules attention to other RB tasks. Currently, this description ignores the cost[26] of switching attention from one task to another.

[26] Cost of attentional switching can be modeled by the time required for the sensory/perceptual observer to converge.

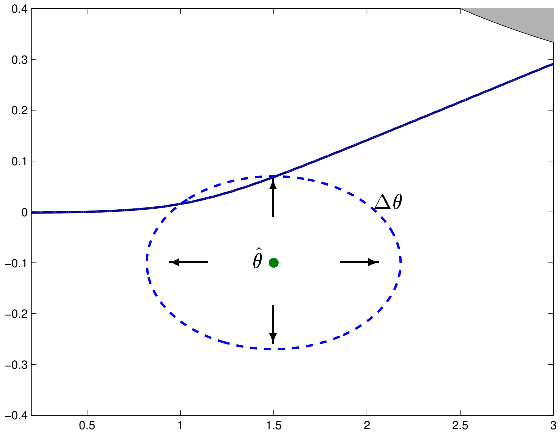

Figure 35. State error growth and attentional updating requirements.

B Brief Tutorial: Phase Plane Trajectories

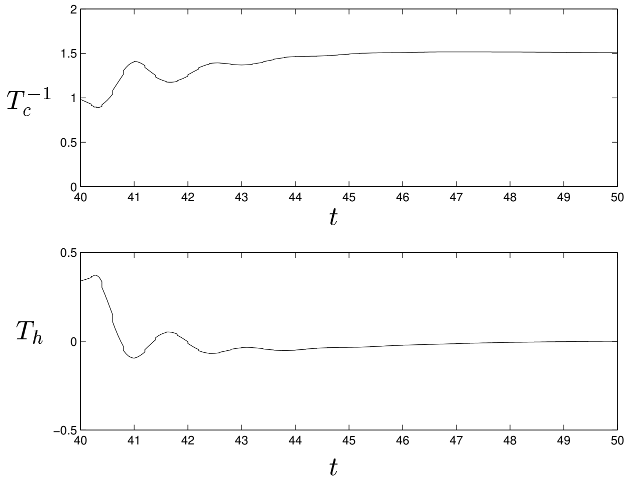

Figure 36. Time histories.

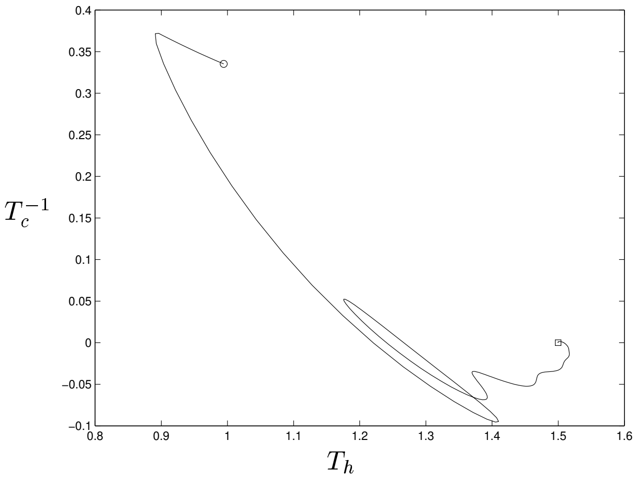

Figure 37. Phase trajectories.

A phase plane trajectory is a way of representing data parameterized by time. For example, suppose that there exist two time series T_h and T_c^{-1}. These time series can be plotted as functions of time as in Figure 36. Though these time histories are useful in understanding the dynamic behavior of the systems that generated them, they do not provide insight into the relations between the two state variables T_h and T_c^{-1}. These relations can be illustrated using the phase plane trajectories wherein T_c^{-1} is plotted against T_h, and where time parameterizes the curve. A phase plane trajectory that corresponds to the time series data in Figure 36 is shown in Figure 37. Times $t = 40$ and $t = 60$ are indicated by $\circ$ and $\Box$, respectively. The oscillatory transients in the time series appear as curves in the the phase trajectories. For this data, the system that generates the phase plane trajectory is driving the system to $T_h = 1.5$ and $T_c^{-1} = 0$, and does so in a spiral curve. We employ these phase plane trajectories throughout the document because they effectively illustrate variable interdependence.

C Summary of Vehicle Dynamics Model Used in CBR Simulator

In the simulations, we employ the nonlinear vehicle dynamics model used in the CBR simulator, which is based upon the vehicle kinetic energy. The vehicle model uses normalized control, $u_k \in [-1, 1] = U_A$, where $u_k < 0$ implies braking, and $u_k > 0$ implies acceleration. We further restrict control authority such that $\Delta u_k \in [-.2, .2] = U_A^\Delta$. The equations of motion are given by

$$\begin{aligned} v_k &= \sqrt{\frac{2E_k}{m}} \\ E_{k+1} &= E_k - v_k \Delta t \left[v_k^2 C_{ad} + C_{sd} m g \right] \\ &\quad + h(v_k; u_k) \\ h(v_k; u_k) &= \begin{cases} \eta \Delta t u_k & u_k \geq 0 \\ \beta \Delta t v_k u_k & u_k < 0 \end{cases}, \end{aligned}$$

where E_k is the vehicle kinetic energy at time k, v_k is the vehicle velocity, C_{ad} is the coefficient of drag due to air resistance, C_{sd} is the coefficient of drag due to surface friction, m is the vehicle mass, g is the acceleration due to gravity, η is the maximum engine output, and β is the maximum brake force. The vehicle parameters are as follows: $m = 1350$ kg, $\eta = 10600$ Nm/s, $\beta = 8000$ N, $C_{ad} = 0.25$ kg/m, $C_{sd} = 0.2$, and $\Delta t = 0.1$ s.

References

[1] J. S. Albus. The Engineering of Mind. Report generated for the Intelligent Systems Division, National Institute of Standards and Technology, Gaithersburg, MD 20895.

[2] J. S. Albus. Outline for a theory of intelligence. *IEEE Transactions on Systems, Man, and Cybernetics*, 21(3):473–509, May/June 1991.

[3] E. R. Boer. Tangent point oriented curve negotiation. In *Proceedings of the 1996 IEEE Intelligent Vehicles Symposium*, pages 7–12, Seikei University, Tokyo, Japan, September 19-20 1996.

[4] E. R. Boer, E. C. Hildreth, and M. A. Goodrich. Drivers in pursuit of perceptual and virtual targets. In *IEEE proceedings of the Intelligent Vehicles '98 Symposium*, Stuttgart, Germany, October 28-30 1998.

[5] R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.

[6] R. A. Brooks. New approaches to robotics. *Science*, 253:1227–1232, September 1991.

[7] Rodney A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2:14–23, 1986.

[8] B. Chaib-draa. Connection between micro and macro aspects of agent modeling. In *Agents'97*, Marina Del Rey, CA, 1997. ACM.

[9] C. L. Connolly and R. A. Grupen. On the applications of harmonic functions to robotics. *Journal of Robotic Systems*, 10(7):931–946, 1993.

[10] J. Dean. Animats and what they can tell us. *Trends in Cognitive Sciences*, 2(2):60–67, February 1998.

[11] D. Driankov, H. Hellendoorn, and M. Reinfrank. *An Introduction to Fuzzy Control.* Springer, 2nd edition, 1996.

[12] A. P. Duchon, W. H. Warren, and L. P. Kaelbling. Ecological robotics. *Adaptive Behavior*, 6(3/4):471–505, 1998.

[13] P. C. Fishburn. Subjective expected utility: A review of normative theories. *Theory and Decision*, 13:139–199, 1981.

[14] G. Gigerenzer and D. G. Goldstein. Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103(4):650–669, 1996.

[15] M. A. Goodrich. *A Theory of Satisficing Control.* Ph.D. Dissertation, Brigham Young University, 1996.

[16] M. A. Goodrich and E. R. Boer. Semiotics and mental models: Modeling automobile driver behavior. In *Joint Conference on Science and Technology of Intelligent Systems ISIC/CIRA/ISAS'98 Proceedings*, Gaithersburg, MD, September 1998.

[17] M. A. Goodrich, R. L. Frost, and W. C. Stirling. A satisficing fuzzy logic controller. To appear in *IEEE Transactions on Fuzzy Systems.*

[18] M. A. Goodrich, W. C. Stirling, and R. L. Frost. A theory of satisficing decisions and control. *IEEE Transactions on Systems, Man, and Cybernetics — Part A: Systems and Humans*, 28(6):763–779, November 1998.

[19] J. Guldner and V. I. Utkin. Sliding Mode Control for an Obstacle Avoidance Strategy Based on an Harmonic Potential Field. In *Proceedings of the 32nd Conference on Decision and Control*, pages 424–429, San Antonio, Texas, December 1993.

[20] J. H. Hogema, W. H. Janssen, M. Coemet, and H. J. Soeteman. Effects of intelligent cruise control on driving behavior — a simulator study. In *ITS 96*, 1996.

[21] P. N. Johnson-Laird. *The Computer and the Mind: An Introduction to Cognitive Science.* Harvard University Press, Cambridge, Massachusetts, 1988.

[22] Jin-Oh Kim and P. K. Khosla. Real-time obstacle avoidance using harmonic potential functions. *IEEE Transactions on Robotics and Automation*, 8(3):338–349, June 1992.

[23] R. Kumar and J. A. Stover. A behavior-based intelligent control architecture. In *Proceedings of the 1998 IEEE ISIC/CIRA/ISAS Joint Conference*, pages 549–553, Gaithersburg, MD, September 14-17 1998.

[24] D. N. Lee. A theory of visual control of braking based on information about time-to-collision. *Perception*, pages 437–459, 1976.

[25] I. Levi. *The Enterprise of Knowledge.* MIT Press, Cambridge, Massachusetts, 1980.

[26] F. L. Lewis. *Optimal Control.* Wiley-Interscience, New York, 1976.

[27] M. J. Matarić. Behavior-based robotics as a tool for synthesis of artificial behavior and analysis of natural behavior. *Trends in Cognitive Sciences*, 2(3):82–87, March 1998.

[28] T. Matsuda and S. Takatsu. Algebraic properties of satisficing decision criterion. *Information Sciences*, 17(3), 221-237 17.

[29] T. Matsuda and S. Takatsu. Characterization of satisficing decision criterion. *Information Sciences*, 17(2):131–151, 1979.

[30] D. Q. Mayne and H. Michalska. Receding horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 35:814–824, 1990.

[31] J. M. Mendel. Fuzzy logic systems for engineering: A tutorial. *Proceedings of the IEEE*, 83(3):345–377, March 1995.

[32] M. D. Mesarovic. Systems theoretic approach to formal theory of problem solving. In R. Banerji and M. D. Mesarovic, editors, *Theoretical Approaches to Non-Numerical Problem Solving*, pages 161–178. Springer, 1970.

[33] M. D. Mesarovic and Y. Takahara. On a qualitative theory of satisfactory control. *Information Sciences*, 4(4):291–313, October 1972.

[34] A. Meystel. Intelligent systems: A semiotic perspective. In *Proceedings of the IEEE International Symposium on Intelligent Control*, Dearborn, Michigan, September 15-18 1996.

[35] H Michalska and D. Q. Mayne. Moving horizon observers and observer-based control. *IEEE Transactions on Automatic Control*, 40(6):995–1006, June 1995.

[36] M. Minsky. *The Society of Mind.* Simon and Schuster, 1986.

[37] Yun Seok Nam, Bum Hee Lee, and Nak Yong Ko. A View-Time Based Potential Field Method for Moving Obstacle Avoidance. *SICE*, pages 1463–1468, 1996.

[38] M. C. Neves and E. Oliveira. A control architecture for an autonomous mobile robot. In *Agents'97*, Marina Del Rey, CA, 1997. ACM.

[39] G. De Nicolao, L. Magni, and R. Scattolini. Stabilizing receding-horizon control of nonlinear time-varying systems. *IEEE Transactions on Automatic Control*, 43(7):1030–1036, July 1998.

[40] J. Rasmussen. Outlines of a hybrid model of the process plant operator. In T. B. Sheridan and G. Johannsen, editors, *Monitoring Behavior and Supervisory Control*, pages 371–383. Plenum, 1976.

[41] J. Rasmussen. Skills, rules, and knowledge; signals, signs, and symbols, and other distinctions in human performance models. *IEEE Transactions on Systems, Man, and Cybernetics*, 13(3):257–266, May/June 1983.

[42] R. Rensink and E. Boer, editors. *Cambridge Basic Research 1996 Annual Report*, Cambridge, Massachusetts, December 1996. Nissan Research and Development, Inc. Technical Report CBR TR 96-10.

[43] J. Richalet. Industrial applications of model based predictive control. *Automatica*, 29:1251–1274, 1993.

[44] G. N. Saridis. Analytic formulation of the principle of increasing precision with decreasing intelligence for intelligent machines. *Automatica*, 25(3):461–467, 1989.

[45] P. O. M. Scokaert, J. B. Rawlings, and E. S. Meadows. Discrete-time stability with perturbations: Application to model predictive control. *Automatica*, 33(3):463–470, March 1997.

[46] T. B. Sheridan. *Telerobotics, Automation, and Human Supervisory Control.* MIT Press, 1992.

[47] H. A. Simon. A behavioral model of rational choice. *Quart. J. Economics*, 59:99–118, 1955.

[48] H. A. Simon. Invariants of human behavior. *Annu. Rev. Psychology*, 41:1–19, 1990.

[49] H. A. Simon. *The Sciences of the Artificial.* MIT Press, 3rd edition, 1996.

[50] P. B. Sistu and B. W. Bequette. Nonlinear model-predictive control: Closed-loop stability analysis. *AIChE Journal*, 42(12):3388–3402, December 1996.

[51] W. C. Stirling and M. A. Goodrich. Satisficing games. *Information Sciences*, 1998. To appear.

[52] S. Sundar and Z. Shiller. Optimal obstacle avoidance based on the Hamilton-Jacobi-Bellman equation. *IEEE Transactions on Robotics and Automation*, 13(2):305–310, April 1997.

[53] S. Takatsu. Decomposition of satisficing decision problems. *Information Sciences*, 22(2):139–148, 1980.

[54] S. Takatsu. Latent satisficing decision criterion. *Information Sciences*, 25(2):145–152, 1981.

[55] S. Wood. A probabilistic network approach to prioritising sensing and reasoning: a step towards satisficing modelling. In *Proceedings of the AAAI Spring Symposium on Satisficing Models*, pages 87–90, Stanford, California, March 23-25 1998. Technical Report SS-98-05.

[56] E. H. Yilmaz and Jr. W. H. Warren. Visual control of braking: A test of the $\dot{\tau}$ hypothesis. *Journal of Experimental Psychology*, 21(5):996–1014, 1995.

[57] C. E. Zsambok and G. Klein, editors. *Naturalistic Decision Making.* Lawrence Erlbaum Associates, 1997.