Probability-based decision making for automated highway driving — Source link ⧉

Axel Niehaus, Robert F. Stengel

Institutions: Princeton University

Published on: 01 Oct 1991 - Vehicle Navigation and Information Systems Conference

Topics: Guidance system and Stochastic modelling

Related papers:

- Probability-based decision making for automated highway driving

- Wiggling through complex traffic: Planning trajectories constrained by predictions

- A multi-agent case-based traffic control scenario evaluation system

- A strategy on situation evaluation for driver assistance systems in commercial vehicles considering pedestrians in urban traffic

- Development and evaluation of a knowledge-based system for traffic congestion management and control

Probability-Based Decision Making for Automated Highway Driving

Axel Niehaus and Robert F. Stengel, *Fellow, IEEE*

Abstract— Real-time, rule-based guidance systems for autonomous vehicles on limited-access highways are investigated. The goal of these systems is to plan trajectories that are safe, while satisfying the driver's requests based on stochastic information about the vehicle state and the surrounding traffic. A rule-based system is used for high-level planning. Given a stochastic model of the traffic situation driven by current measurements, the probable evolution of traffic and the best trajectory to follow are predicted. Simulation results assess the impact of uncertain knowledge of traffic on the performance of the guidance system, showing that uncertainty can and must be taken into account.

I. INTRODUCTION

INTELLIGENT Vehicle/Highway Systems (IVHS) have been proposed to help reduce the nation's highway congestion problem and to increase the safety of our vehicles [1], [2]. Many discussions have taken place about the use of the term "intelligent" for such systems; the meaning of intelligence here is simply that IVHS systems perform tasks that were previously performed by humans, making use of advanced technologies such as image processing, automatic control, or artificial intelligence.

Depending on where the "intelligence" is placed, a given IVHS system belongs to one of three categories: Intelligent Vehicle, Intelligent Highway, or Intelligent Vehicle *and* Highway. Functions of the first category include vehicle platooning, obstacle avoidance, and autonomous vehicle guidance [3]–[8]. Examples of Intelligent Highway functions include automated toll booths, adaptive traffic signals, and centralized traffic control [9]–[12]. Dynamic navigation, where traffic density information is used by the vehicle to select a path, is an example of the third category, requiring an intelligent environment as well as on-board intelligence.

An intelligent guidance system for an autonomous vehicle operating on a limited-access highway is presented here, building on the Intelligent Guidance for Headway and Lane Control (IGHLC) System presented in previous papers [6]–[8]. Fig. 1 shows how this guidance system is integrated with the remaining systems to perform closed-loop control of the
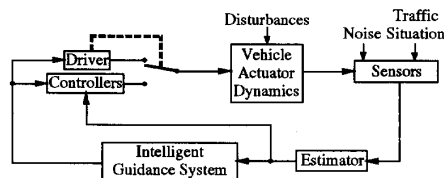
Fig. 1. Closed-loop control of an autonomous vehicle.

vehicle. The guidance task consists of an analysis of the information received from estimators to plan alternative trajectories for the autonomous vehicle. No communication of intended actions is assumed between vehicles. Radar, magnetometers combined with specially prepared highways, and computer vision have been proposed as sources of guidance information [13]–[16]. These technologies provide varying degrees of precision at a variety of costs, but none yields error-free measurements.

A method is proposed for handling uncertain information explicitly. The associated analysis reveals tolerable levels of uncertainty, which is useful information for evaluating competing sensor systems. Probability theory is used to describe uncertainty, providing a solid and rigorous background while being compatible with the intuitive notion of the likelihood [17]–[19]. Applying standard probabilistic methods to automated highway driving is made difficult by the real-time nature and complexity of the problem. Techniques have been developed to make the use of probability theory for autonomous-vehicle guidance more manageable, and they are applied here [8].

The first section of the paper introduces IGHLC, and establishes the need for a method to predict how traffic situations may evolve under uncertainty. The second section presents a probabilistic framework for automated decision making. The third section shows how this framework is applied to IGHLC. The fourth section provides simulation results that illustrate the guidance system's performance in a sample traffic situation with different levels of uncertainty. The final section summarizes the benefits and limitations of the proposed methodology.

II. INTELLIGENT GUIDANCE FOR HEADWAY AND LANE CONTROL

IGHLC is a rule-based expert system that performs the guidance function for an autonomous vehicle on a limited-access highway. The inputs to the system include the coordinates and
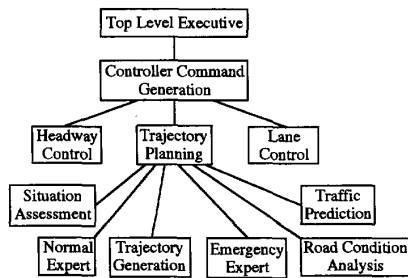
Fig. 2. Functional decomposition of IGHLC knowledge base.

velocities of the autonomous ("own") vehicle and surrounding traffic, the road geometry, the current road condition (e.g., ice, snow), and driver-selectable parameters consisting of a target cruising speed and desired levels of safety and aggressiveness. At every iteration of the guidance loop, the task of the expert system is to analyze this information and to determine an appropriate trajectory for the Own Vehicle, resulting in a lateral command (e.g., change lane or continue straight) and a longitudinal command (e.g., accelerate or achieve target velocity). These commands are implemented by controllers for "inner-loop" control of the vehicle (Fig. 1). The remainder of this section provides a brief overview of the IGHLC expert system. The reader interested in more details (e.g., motivation of the expert system approach, rule base composition, inference engine) is referred to [7].

The expert system's knowledge base is composed of several modules, each responsible for a specific function (Fig. 2). The task of the Situation Assessment function is to determine whether the situation is safe or unsafe for the Own Vehicle. The safety of the relative position of a vehicle A_n and the vehicle A_{n+1} in front of it is defined in terms of an *individual time ratio* [8]

$$t_R(A_n, A_{n+1}) = \frac{\delta_n}{t_d v_n} \qquad (1)$$

and an *individual braking ratio*

$$b_R(A_n, A_{n+1}) = \begin{cases} \dfrac{(v_{n+1} - v_n)^2}{2\, d_n \delta_n} & \text{if } v_{n+1} < v_n \\ 0 & \text{if } v_{n+1} \geq v_n \end{cases} \qquad (2)$$

where v_i designates the velocity of A_i, δ_i is the distance between A_i and A_{i+1}, t_d is the desired separation time between vehicles, and d_i is the maximum obtainable deceleration of A_i, calculated by the Road Condition Analysis function. If A_{n+1} is in front of A_n and A_{n-1} is behind, two *global* ratios are defined as

$$t_G(A_n) = \min\left[t_R(A_{n-1}, A_n), t_R(A_n, A_{n+1})\right] \qquad (3)$$

and

$$b_G(A_n) = \max\left[b_R(A_{n-1}, A_n), b_R(A_n, A_{n+1})\right]. \qquad (4)$$

A_n is considered safe if $t_G(A_n) \geq 1/2$ and $b_G(A_n) \leq 1/2$. A collision is likely if $t_G(A_n) < 0$ or $b_G(A_n) > 1$, and an emergency is declared in all other cases.

In both safe and unsafe situations, the expert system must identify an adequate acceleration to use and must determine if lane changes are desired. The Trajectory Generation function is used to predict the trajectories of the Own Vehicle and surrounding traffic for a set of discrete levels of acceleration that may be used by the Own Vehicle. Once the candidate trajectories have been obtained, the expert system selects the best one.

In unsafe conditions, the Emergency Expert selects the option that maximizes safety. A trajectory leading to an accident is always considered worse than one that avoids collisions. If several options avoid collisions, the trajectory Γ minimizing the cost function

$$J_{\text{tot}}(\Gamma) = c_t \max_{\text{on}\Gamma} (J_t) + c_b b_G^f \qquad (5)$$

is selected, where c_t and c_b are real coefficients, J_t is the cost associated with the global time ratio of the Own Vehicle at a given point on Γ, and b_G^f is the global braking ratio of the Own Vehicle at the final position on the trajectory

$$J_t = \begin{cases} 0 & \text{if } t_G(\text{Own}) \geq 1 \\ 1 - t_G(\text{Own}) & \text{if } t_G(\text{Own}) < 1. \end{cases} \qquad (6)$$

In safe situations, the Normal Expert selects an option that not only is safe, but that satisfies the driver's goal. If obstacles (including slower vehicles) are detected ahead in the same lane or in a lane to the left of the Own Vehicle, the Normal Expert selects a safe left-lane change trajectory (if available) for passing. Otherwise, the Normal Expert guides the Own Vehicle into the rightmost lane; if that lane contains obstacles, they are passed before changing the lane.

The performance of IGHLC depends largely on the accuracy of the predictions made by the Trajectory Generation function. In our past work, these predictions were limited to the longitudinal motions of the surrounding traffic (no lane changes were predicted for surrounding vehicles), and they did not take into account uncertainty. The next section presents a method for predicting lateral and longitudinal motions in probabilistically defined traffic situations.

III. WORST-CASE DECISION MAKING

Worst-Case Decision Making (WCDM) predicts the evolution of a controlled dynamic system's state and its dynamic environment, assuming that the environment may react in several ways to control histories. It selects the *worst plausible evolution* as a basis for allocating resources.

Consider the control of the discrete-time dynamic system defined by

$$\boldsymbol{x}_{k+1} = \boldsymbol{f}(\boldsymbol{x}_k, k, \boldsymbol{u}_k) + \boldsymbol{\nu}_k \qquad (7)$$

where $\boldsymbol{x}_k \in \mathbb{R}^n, \boldsymbol{u}_k \in \mathbb{R}^n$, and $\boldsymbol{\nu}_k \in \mathbb{R}^n$ are the state, control, and disturbance vectors at time index k. This system is surrounded by an environment modeled by

$$\boldsymbol{y}_{k+1} = \boldsymbol{g}(\boldsymbol{y}_k, k, \boldsymbol{v}_k, \boldsymbol{w}_k) + \boldsymbol{\eta}_k \qquad (8)$$

where $\boldsymbol{y}_k \in \mathbb{R}^p$ and $\boldsymbol{\eta}_k \in \mathbb{R}^p$ are the environment state and disturbance vectors. Exterior agents are assumed to control

the environment through the control vectors $v_k \in \mathbb{N}_S^{*A}$ and[1] $w_k \in \mathbb{R}^r$. The former represents a set of S distinct strategies that may be chosen by A agents. The latter represents a set of r additional continuous controls. The state vectors x_k and y_k are assumed to be correlated Gaussian random variables prescribed by their mean values and covariance matrices

$$\hat{x}_k = E(x_k) \tag{9}$$

$$P_k = E[(x_k - \hat{x}_k)(x_k - \hat{x}_k)^T] \tag{10}$$

$$\hat{y}_k = E(y_k) \tag{11}$$

$$Q_k = E[(y_k - \hat{y}_k)(y_k - \hat{y}_k)^T] \tag{12}$$

$$R_k = E[(x_k - \hat{x}_k)(y_k - \hat{y}_k)^T] \tag{13}$$

where $E[\cdot]$ is the expectation operator. The disturbance inputs ν_k and η_k are two uncorrelated zero-mean white Gaussian random sequences

$$E(\nu_k) = 0 \tag{14}$$

$$E(\nu_k \nu_k^T) = \Lambda_k \tag{15}$$

$$E(\eta_k) = 0 \tag{16}$$

$$E(\eta_k \eta_k^T) = M_k. \tag{17}$$

The goal of WCDM is to predict how x_k and y_k evolve when a given control history u_k is used. As in the extended Kalman filter [20], the mean $\hat{x}_k$ is propagated by

$$\hat{x}_{k+1} = f(\hat{x}_k, k, u_k). \tag{18}$$

The covariance P_k is propagated using a linearization of f about the mean value

$$P_{k+1} = \Lambda_k + \Phi_k P_k \Phi_k^T \tag{19}$$

where

$$\Phi_k = \frac{\partial f}{\partial x}(\hat{x}_k, k, u_k). \tag{20}$$

The exterior agents must be modeled to propagate y_k, Q_k, and R_k. The power of WCDM lies in the fact that it allows multiple strategies for the agents to be compared. Each of the A components of v_k can take integer values between 1 and S, modeling A agents, each having the choice between S strategies (which may vary over the agents). WCDM assumes that a model is available to predict which of the S strategies are plausible for each agent. Depending on whether or not strategy j is judged plausible for agent i at a time step k, the value of the *plausibility function* λ is ± 1

$$\lambda(i, j, x_k, y_k, k) = \begin{cases} +1 & \text{if } [v_k(i) = j] \text{ plausible} \\ -1 & \text{otherwise.} \end{cases} \tag{21}$$

Furthermore, WCDM assumes that the selection of the continuous controls w_k by the agents can be modeled in a more standard fashion by

$$w_k = \mu(x_k, y_k, k, v_k) \tag{22}$$

essentially stating that the continuous controls w_k can be predicted in a more precise manner once the strategies v_k are known.

[1] $\mathbb{N}_S^* = \{1, \cdots, S\}$.

In a first step, WCDM determines the set of plausible strategies for the external agents, taking into account the probability distributions of x_k and y_k. The *belief interval function* b associates a finite interval of real numbers to the mean m and variance v of a random variable according to

$$b(m, v) = [m - \kappa\sqrt{v}, m + \kappa\sqrt{v}] \tag{23}$$

where κ is a real positive constant, bounding the distribution of the random variable to κ standard deviations on each side of the mean. The *expected plausibility function* $\hat{\lambda}$ is defined by

$$\hat{\lambda}(i, j, \hat{x}_k, P_k, \hat{y}_k, Q_k, k) = \begin{cases} +1 & \text{if } \mathcal{P} \text{ is } true \\ -1 & \text{otherwise} \end{cases} \tag{24}$$

where the proposition $\mathcal{P}$ is given by

$$\mathcal{P} = \left\{ \begin{array}{c} \text{There exist } x^* \in \mathbb{R}^n \text{ and } y^* \in \mathbb{R}^p \\ \text{such that for all } s \in \mathbb{N}_n^* \text{ and } t \in \mathbb{N}_p^* \\ \left\{ \begin{array}{l} x^*(s) \in b[\hat{x}_k(s), P_k(s, s)] \\ y^*(t) \in b[\hat{y}_k(t), Q_k(t, t)] \\ \hat{\lambda}(i, j, x^*, y^*, k) = 1 \end{array} \right. \end{array} \right\}. \tag{25}$$

The arguments of x, y, P, and Q designate components of the corresponding vectors and matrices. $\hat{\lambda}$ indicates whether or not values of x^* and y^* exist close enough to the means $\hat{x}_k$ and $\hat{y}_k$ to make strategy j plausible for agent i.

In a second step, WCDM identifies the strategy for each agent among the set of strategies (judged plausible in the first step) that would be the worst for the dynamic system. The hazard presented by strategy j of agent i is assumed to be given by the *expected hazard function* $\hat{h}$ where

$$\hat{h}(i, j, \hat{x}_k, P_k, \hat{y}_k, Q_k, k) \in \mathbb{R}^+. \tag{26}$$

The worst plausible strategy of agent i is then used as the estimate for $v_k(i)$

$$\hat{v}_k(i) = \arg\max_{j \in \mathbb{N}_S^*} \hat{\lambda}[1 + \hat{h}]. \tag{27}$$

Having obtained $\hat{v}_k$, the mean $\hat{y}_k$ is propagated by

$$\hat{y}_{k+1} = g(\hat{y}_k, k, \hat{v}_k, \mu(\hat{x}_k, \hat{y}_k, k, \hat{v}_k)). \tag{28}$$

The covariances Q_k and R_k are propagated using linearizations of g and μ about the mean values

$$Q_{k+1} = M_k + A_k[Q_k A_k^T + R_k^T B_k^T]$$
$$+ B_k[R_k A_k^T + P_k B_k^T] \tag{29}$$

$$R_{k+1} = \Phi_k[R_k A_k^T + P_k B_k^T] \tag{30}$$

where

$$\Psi_k^y = \frac{\partial g}{\partial y}[\hat{y}_k, k, \hat{v}_k, \mu(\hat{x}_k, \hat{y}_k, k, \hat{v}_k)] \tag{31}$$

$$\Psi_k^w = \frac{\partial g}{\partial w}[\hat{y}_k, k, \hat{v}_k, \mu(\hat{x}_k, \hat{y}_k, k, \hat{v}_k)] \tag{32}$$

$$\Gamma_k^x = \frac{\partial \mu}{\partial x}(\hat{x}_k, \hat{y}_k, k, \hat{v}_k) \tag{33}$$

$$\Gamma_k^y = \frac{\partial \mu}{\partial y}(\hat{x}_k, \hat{y}_k, k, \hat{v}_k) \tag{34}$$

$$A_k = \Psi_k^y + \Psi_k^w \Gamma_k^y \tag{35}$$

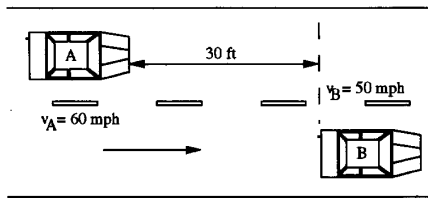$$B_k = \Psi_k^w \Gamma_k^x. \tag{36}$$

Fig. 3. Sample traffic situation for worst-case demonstration.

Equations (18)–(20), (23)–(25), and (27)–(36) enable the prediction of the worst plausible evolution of the system and evnironmental states. At each iteration, estimates are obtained for the states, along with corresponding covariances. An important application of the method, as presented in the following section, is the determination of an optimal control u_k^* by minimization of a stochastic cost function defined on the predicted evolutions.

IV. APPLICATION OF WORST-CASE DECISION MAKING TO AUTOMOTIVE GUIDANCE

As shown in the first section, the Trajectory Generation Function of the IGHLC expert system must predict how the current traffic situation will evolve under a given control strategy for the Own Vehicle. This section first examines why WCDM is well suited to this problem, then presents the implementation, and finally shows how the predicted evolutions are used by the expert system.

A. Why Use Worst-Case Decision Making?

Consider the traffic situation shown in Fig. 3, where Vehicle A is about to pass Vehicle B. Under the assumptions made in IGHLC, Vehicle B's intended action is not available to Vehicle A's IGHLC system. The latter must decide whether or not Vehicle B will suddenly change to the left based on its own observations (i.e., state of turn signals, coordinates, and velocities). In the absence of any additional information, it appears reasonable to assume that Vehicle B will not perform the lane change, and the passing maneuver would continue. Assume now that Vehicle A's IGHLC system has made an observation suggesting that Vehicle B will perform the lane change with a probability of 0.2. In this case, full braking might be applied to avoid the potential collision, even though this option is four times less likely to occur than the other one.

As seen from this example, events may be considered possible only when their probabilities are higher than a threshold. If several events are judged possible, it might be preferred to plan for the most dangerous one, even if that is not the most probable event. This attitude is adopted by WCDM.

B. Applying Worst-Case Decision Making to Highway Driving

State Transition Functions f and g: For IGHLC, the n-dimensional system state x_k corresponds to the Own Vehicle. Its components are the longitudinal road coordinate of the center of the Own Vehicle, the longitudinal velocity of the Own Vehicle, and the lane number of the Own Vehicle (which is constant for the current application). The control u_k is

simply the longitudinal acceleration of the Own Vehicle

$$x_{k+1} = \begin{bmatrix} x_{k+1}^o \\ v_{k+1}^o \\ \epsilon_{k+1}^o \end{bmatrix} = \begin{bmatrix} 1 & \Delta t & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_k^o \\ v_k^o \\ \epsilon_k^o \end{bmatrix} + \begin{bmatrix} \frac{\Delta t^2}{2} \\ \Delta t \\ 0 \end{bmatrix} u_k + \nu_k. \tag{37}$$

The environment consists of A vehicles, the drivers of these vehicles constituting the A exterior agents of WCDM. Each vehicle is characterized by the same variables as the Own Vehicle, and the continuous control vector w_k consists of the A longitudinal accelerations $(r = A)$. The discrete control vector ν_k corresponds to the lateral decisions made by the A drivers: +1 indicates a right-lane change, 0 corresponds to no-lane change, and −1 indicates a left-lane change. The equation governing the evolution of the state of Vehicle i is

$$y_{k+1}(i) = \begin{bmatrix} x_{k+1}^i \\ v_{k+1}^i \\ \epsilon_{k+1}^i \end{bmatrix}$$
$$= \begin{bmatrix} 1 & \Delta t & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_k^i \\ v_k^i \\ \epsilon_k^i \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$
$$\cdot v_k(i) + \begin{bmatrix} \frac{\Delta t^2}{2} \\ \Delta t \\ 0 \end{bmatrix} w_k(i) + \eta_k(i). \tag{38}$$

Plausibility Function λ: The plausibility function λ specifies which lateral actions are presumed plausible for each vehicle driver. The system presented here computes three scores for each vehicle at every step: one for a left-lane change, one for a right-lane change, and one for continuing straight. Each score reflects the estimated likelihood of the corresponding lateral option. For every vehicle, all three scores are initially set to zero and then modified according to Table I. The following paragraph explains how Table I is used for a given vehicle (among the A surrounding vehicles).

The first two lines of the table account for the existence of lanes to the left and right of the considered vehicle. The scores are then modified to reflect the safety of each lane and to account for eventual obstacles (lines 3–10). The score of an option leading to a collision is decreased by 10 (which eliminates the option); if the option leads only to an emergency, the score is decremented by one. Safe options are not affected. If the lane in front of the vehicle contains an obstacle but is judged safe, the right-lane score is decreased by one, since this would motivate a left-lane change. Conversely, if the lane does not contain any obstacles and is judged safe, the same-lane score is increased by one, as no reason for a left-lane change exists. The right-lane score also would be increased by one if similarly free of obstacles and safe. This scoring reflects the fact that vehicles should travel in the rightmost lane when not performing a passing maneuver. The last four lines of Table I show how the scores are affected by direct observation of the considered vehicle. A strong left-lane change indication, such as observation of the left-turn

TABLE I
COMPUTATION OF LANE SCORES FOR
DETERMINATION OF PLAUSIBLE LATERAL ACTIONS

Observation	Left-Lane Score Increment	Same-Lane Score Increment	Right-Lane Score Increment
Left Lane Doesn't Exist	-100	0	0
Right Lane Doesn't Exist	0	0	-100
Same Lane Safe Without Obstacle	0	1	0
Same Lane Safe With Obstacle	0	0	-1
Same Lane Emergency	0	-1	0
Left Lane Emergency	-1	0	0
Left Lane Collision	-10	0	0
Right Lane Emergency	0	0	-1
Right Lane Collision	0	0	-10
Same & Right Lane Safe Without Obstacles	0	0	1
Strong Left-Lane Change Indication	2	0	-2
Weak Left-Lane Change Indication	1	0	-1
Weak Right-Lane Change Indication	-1	0	1
Strong Right-Lane Change Indication	-2	0	2

signal and high lateral speed to the left, increases the left-lane score by two while decreasing the right-lane score by the same amount. Finally, if all three options correspond to emergencies but one is significantly safer than the others, that option receives an additional point.

The plausibility function is set to one for the options obtaining the highest score. This may occur for several options, reflecting the impossibility of precisely predicting the lateral actions of a human driver. λ is set to -1 for the remaining options, eliminating those with insufficient credibility. If all vehicles were operating with IGHLC, a more precise model could be used for lateral predictions.

Expected Hazard Function $\hat{h}$: For the IGHLC implementation, it is assumed that the most hazardous lateral action of a surrounding vehicle is to enter (or stay in) the Own Vehicle's lane, as long as that would place it directly in front of or behind the Own Vehicle. In all other cases, the most hazardous action is set to be no lane change as a default.

Model for Continuous Controls μ: The function μ (22) must predict the accelerations of the surrounding vehicles in terms of the traffic situation and the estimated lateral actions (27). Given the velocities v_A of a Vehicle A and v_B of the Vehicle B that would be ahead of it after execution of A's estimated lateral action, the separation δ between A and B, and A's predicted desired velocity α, μ predicts A's acceleration as follows.

• *If Vehicle B exists*, the maximum allowed acceleration for minimum safety margins is computed

$$\bar{a} = \min\left(\bar{a}_1, \bar{a}_2\right). \tag{39}$$

$\bar{a}_1$ limits the maximum allowed acceleration in terms of the distance δ between the vehicles

$$\bar{a}_1 = \begin{cases} \dfrac{10a[\delta - t_d(1 + s_F)v_A]}{t_d(1 + s_F)v_A(5 - 4s_F)} & \text{if } \dfrac{\delta}{t_d v_A} \geq 1 + s_F \\ \dfrac{10d[\delta - t_d(1 + s_F)v_A]}{t_d(1 + s_F)v_A(5 - 4s_F)} & \text{otherwise} \end{cases} \tag{40}$$

where a and d are the assumed maximum obtainable levels of acceleration and deceleration, and s_F is A's presumed *security factor* (varying between 0 and 1).

$\bar{a}_2$ limits the maximum allowable acceleration in terms of the deceleration $\bar{d}$ required to avoid a collision with B

$$\bar{a}_2 = \begin{cases} a\left(\dfrac{d - 10\bar{d}}{d}\right) & \text{if } \bar{d}/d \leq 0.1 \\ \min\left(3, \dfrac{10\bar{d} - d}{d}\right)\bar{d}\dfrac{10.5 + s_F}{-30} & \text{if } \bar{d}/d > 0.1 \end{cases} \tag{41}$$

where

$$\bar{d} = \begin{cases} \dfrac{(v_A - v_B)^2}{2\delta} & \text{if } v_B < v_A \\ 0 & \text{if } v_B \geq v_A. \end{cases} \tag{42}$$

Equation (41) limits the allowed acceleration if a small deceleration is necessary to avoid $B(\bar{d} \leq 0.1\,d)$, and it forces a certain percentage of $\bar{d}(0\%$ if $\bar{d} = 0.1d$ up to more than 100% for $\bar{d} \geq 0.4\,d)$ to be used when $\bar{d}$ becomes larger. The system proceeds by comparing v_B to the desired velocity α.

–If $v_B < \alpha$, then B is an obstacle, and A is assumed to regulate its velocity to v_B while keeping a safe distance. The gross controller output

$$a_{\text{Gross}} = 0.0447[\delta - t_d(1 + s_F)v_B] + 0.304[v_B - v_A] \tag{43}$$

and μ predicts A's acceleration to be

$$\hat{a} = \max\left\{-d, \min\left[\bar{a}, a(0.4 + 0.6a_F), a_{\text{Gross}}\right]\right\} \tag{44}$$

where a_F is A's presumed *aggressiveness factor* (varying between 0 and 1).

–If $v_B \geq \alpha$, then B is not an obstacle, and μ assumes that A regulates its velocity to the desired value

$$\hat{a} = \max\left\{-d, \min\left[\bar{a}, a(0.4 + 0.6a_F), 0.2(\alpha - v_A)\right]\right\}. \tag{45}$$

• *If Vehicle B does not exist*, the system assumes that A regulates its velocity to α. In this case, the predicted acceleration for A is given by

$$\hat{a} = \begin{cases} \min\left[\dfrac{a(4 + 6a_F)}{10}, \dfrac{\alpha - v_A}{5}\right] & \text{if } v_A < \alpha \\ \max\left[\dfrac{-d(2 + 3a_F)}{10}, \dfrac{\alpha - v_A}{5}\right] & \text{if } v_A \geq \alpha. \end{cases} \tag{46}$$

C. Interfacing Worst-Case Decision Making with the IGHLC Expert System

The task of the Situation Assessment function of IGHLC is to determine the safety of the current traffic situation and to relay control to either the Normal or Emergency Expert. With WCDM, the expert system transfers control to the Normal Expert, as if the situation were known to be safe. In the process
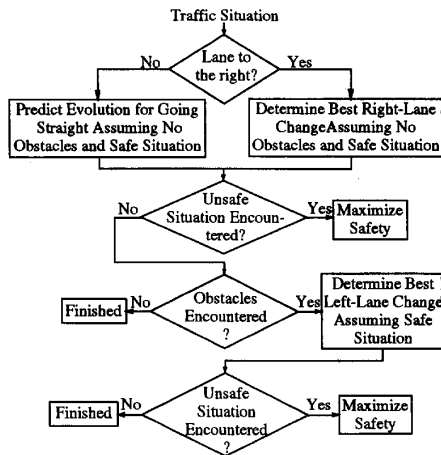
Fig. 4. Normal expert logic.

Fig. 5. Initial traffic situation for guidance system demonstration.

trajectory is first related to the probability of a collision for the Own Vehicle.

$$\Pr\left[\text{Collision}\left(t_k\right)\right]$$

$$= \begin{cases} \Pr\left\{t_G\left[\text{Own}\left(t_k\right)\right] < 0\right\} & \text{if } t_k \neq t_f \\ \Pr\left\{\begin{array}{c} t_G[\text{Own}(t_k)] < 0 \\ \text{or} \\ b_G[\text{Own}(t_k)] > 1 \end{array}\right\} & \text{if } t_k = t_f \end{cases} \tag{48}$$

where t_f is the last predicted instant for the trajectory. The special case $t_k = t_f$ in (48) is used to account for trajectories that have been predicted up to points just prior to dangerous situations, taking into account the global braking ratio of the Own Vehicle at the final positions.

If several trajectories have most critical points with negligible collision probabilities, the Emergency Expert selects the option corresponding to the trajectory that minimizes the cost function

$$J_{\text{tot}}(\Gamma) = c_t \max_{\text{on}\Gamma}\left[E(J_t)\right] + c_b E(b_G^f) \tag{49}$$

where J_t and b_G^f are the deterministic global time ratio cost and global braking ratio of the Own Vehicle, defined in the first section of the paper. A detailed description of how the expected values and probabilities involved in (48) and (49) are computed can be found in [8].

V. SIMULATION RESULTS

The IGHLC expert system has been implemented in Common LISP on a NeXT computer, with numerical WCDM procedures written in C. This section provides a step-by-step description of the analysis performed by the guidance system when presented with a sample traffic situation. Both low-uncertainty and high-uncertainty scenarios are considered, showing the corresponding effect on the action chosen by the guidance system for the Own Vehicle. Fig. 5 shows a top view of the sample traffic situation where the Own Vehicle (dark gray) is surrounded by three vehicles (white). The parameters corresponding to this situation are given in Table II, and include x (expected value of longitudinal coordinate of vehicle center), v (expected value of velocity), a (maximum obtainable acceleration), d (maximum obtainable deceleration), a_F (aggressiveness factor), s_F (security factor), t_d (desired separation time), α (desired velocity), and l (length of vehicle).

For the low-uncertainty scenario, the standard deviations are chosen according to Table III, σ_x corresponding to the standard deviation in the longitudinal coordinate and σ_v to the standard deviation in the velocity. Since a lane exists to the right of the

of selecting an appropriate action, the Normal Expert uses[2] WCDM to predict the evolution of the situation, revealing its safety. If the situation is predicted to become unsafe, control is passed to the Emergency Expert for maximum safety.

The deterministic safety definitions given by (1)–(4) are generalized to traffic situations prescribed by probability distributions. A situation is considered safe for Vehicle A_n if

$$\Pr\left\{\begin{array}{c} t_G(A_n) \geq 1/2 \\ \text{and} \\ b_G(A_n) \leq 1/2 \end{array}\right\} > threshold \tag{47}$$

where $threshold \in [0, 1]$ is usually chosen close to one.

Normal Expert: The Normal Expert first uses WCDM to obtain an appropriate longitudinal command to guide the Own Vehicle into the rightmost lane, assuming that no obstacles will be encountered. Depending on whether a lane exists to the right of the Own Vehicle, this may require staying in lane or performing a right-lane change (Fig. 4). If no unsafe situations or obstacles are detected, the search has concluded. If obstacles are encountered, the Normal Expert uses WCDM to obtain an appropriate longitudinal command for a left-lane change to pass them (assuming there is a lane to the left of the Own Vehicle). If an unsafe situation is detected, the Normal Expert transfers control to the Emergency Expert to maximize safety. Left-lane change trajectories are analyzed similarly for safety (Fig. 4).

Emergency Expert: The task of the Emergency Expert is to find the safest option for the Own Vehicle. WCDM predicts trajectories for the Own Vehicle and surrounding vehicles, assuming different accelerations of the Own Vehicle. The Emergency Expert selects the acceleration and lateral command that yield the safest predicted trajectory for the Own Vehicle.

Trajectories are compared on the basis of the most critical instant along them. The safety at a given instant t_k along a

[2] The implementation of WCDM consists of a group of functions written in C on a NeXT computer. Calls to these functions are included in the actions of the rules requiring the WCDM procedures, parameters being passed between LISP and C using the interface between the two languages on the NeXT computer.

TABLE II
PARAMETERS OF TRAFFIC SITUATION SHOWN IN FIG. 5

Vehicle	Own	A	B	C
x (ft)	0	10	30	130
v (ft/s)	90	100	70	65
a (ft/s^2)	10	10	10	10
d (ft/s^2)	10	10	10	10
a_F	0.5	0.5	0.5	0.5
s_F	0.5	0.5	0.5	0.5
t_d (s)	2	2	2	2
α (ft/s)	100	100	100	65
l (ft)	13.44	13.44	13.44	13.44

TABLE III
STANDARD DEVIATIONS IN LOW-UNCERTAINTY SCENARIO

Vehicle	A	B	C
σ_x (ft)	0.5	1	3
σ_v (ft/s)	1	1	3

TABLE IV
INITIAL SAFETY EVALUATIONS IN LOW-UNCERTAINTY SCENARIO

Couple	t_R	b_R	Safety
Worst B/C	0.56	0.09	Safe
Best Own/B	0.10	0.92	Emer.
Best Own/A	-0.01	0	Coll.

Own Vehicle, the guidance system starts by examining right lane changes (Fig. 4). The relative position of Vehicles B and C is known to be safe (Table IV), and Vehicle C is predicted to remain in lane. Vehicle B is similarly predicted to remain in its lane: a left-lane change would engender an emergency or a collision with the Own Vehicle (Table IV). The same applies for Vehicle A (replacing left-lane change by right-lane change).

When WCDM is used to predict the impact of a constant velocity trajectory for the Own Vehicle, Vehicle B is predicted to change to the left at 4 s (Fig. 6(a)), and a right-lane change is predicted for Vehicle A at 8.4 s (Fig. 6(b)). Since no safety problems or obstacles are encountered, it is concluded that the best option for the Own Vehicle is to accelerate to the desired speed and to perform a right-lane change as soon as a safe distance ahead of Vehicle C is reached. The rest of this section shows how this decision would be altered by a large uncertainty in Vehicle C's velocity.

The high-uncertainty scenario is identical to the low-uncertainty one, except that the standard deviation in Vehicle C's velocity is increased to 15 ft/s. When examining right-lane changes for the Own Vehicle, the expert system still predicts that Vehicle B will remain in its lane (the minimum cost for a left-lane change is more than 10% higher than the maximum cost for staying in lane) (Table V). However, Vehicle C is immediately predicted to change lanes to the left (Fig. 7(a)) because the left-lane change can yield a safer situation than remaining in lane (Table V).
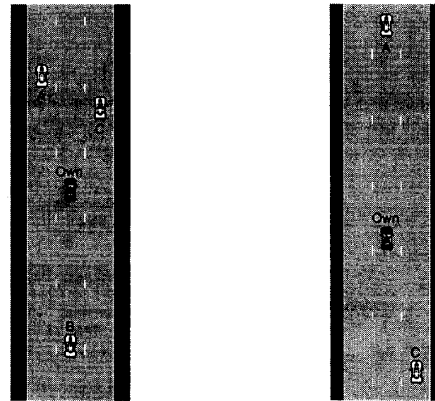

(a)   (b)

Fig. 6. Traffic situations in low-uncertainty scenario. (a) t = 4 s. (b) t = 8.4 s.

TABLE V
INITIAL SAFETY EVALUATIONS IN HIGH-UNCERTAINTY SCENARIO

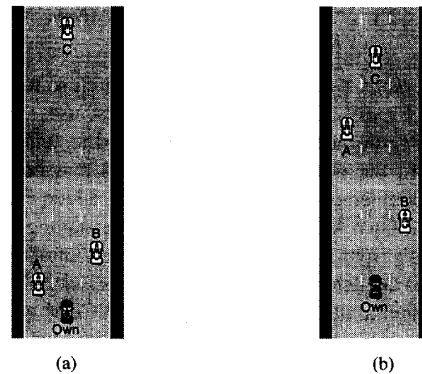| Couple	t_R	b_R	Cost	Safety
Worst B/C	0.56	0.63	1.07	Emer.
Best Own/B	0.10	0.92	1.82	Emer.
Best Own/C	0.68	0.00	0.32	Safe

(a)   (b)

Fig. 7. Traffic situations in high-uncertainty scenario. (a) t = 0.2 s. (b) t = 2.6 s.

Vehicle C is now in the middle lane, causing a possibly dangerous situation for the Own Vehicle. The right-lane change procedure is aborted, and the guidance system begins to maximize safety (Fig. 4). The guidance system's new task is to find the acceleration (or deceleration) that yields the safest left-lane change, right-lane change, or straight trajectory. Fig. 8 shows plots of the collision probability and expected cost for the three options, the acceleration varying between ± 10 ft/s^2. The safest trajectory is obtained with a left-lane change and *maximum* deceleration, yielding a negligible collision probability and an expected cost of 0.56. It is important to keep in mind, though, that these situations represent the prediction of the *worst plausible* evolution; in reality, Vehicle C may well continue in its lane.
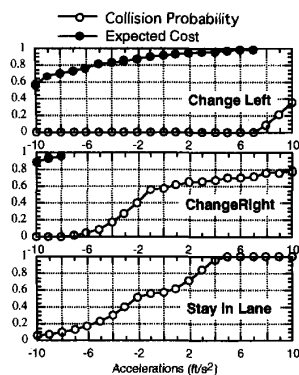
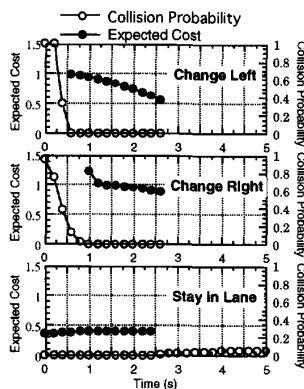Fig. 8. Safety versus acceleration in high-uncertainty scenario.

Fig. 9. Evolution of safety for maximum deceleration in high-uncertainty scenario.

To determine the safest lateral command for the Own Vehicle, plots are generated for the evolution of the safety for the three options when maximum deceleration is used (Fig. 9). The bottom plot shows the evolution of the safety of the *worst* instant for the trajectory in the middle lane. This safety initially decreases as the Own Vehicle approaches Vehicle C, then stabilizes at a collision probability of 0.049 when the distance between the Own Vehicle and Vehicle C increases and the worst instant has been obtained. The shortest distance between the Own Vehicle and Vehicle C is obtained at 2.6 s, and its value is 101.3 ft (Fig. 7(b)). Left-lane changes are initially extremely unsafe (collision probability = 1) and gradually become better as Vehicle A moves away (Fig. 9, top plot).

Left-lane changes are not computed after 2.8 s since left-lane-change trajectories after that time would include the nonzero collision probability of the straight trajectory. The safest left-lane change takes place at 2.6 s. Right-lane changes evolve in a similar way, gradually becoming better as the Own Vehicle passes Vehicle B. The best right-lane change also takes place at 2.6 s, but its expected cost is 0.88. The safest action for the Own Vehicle is thus to stay in lane, using maximum deceleration, then perform a left-lane change at $t = 2.6$ s.

Altering the standard deviation of Vehicle C's velocity has completely changed the predicted worst-plausible evolution of the traffic situation. Rather than accelerating for a future right-lane change, the analysis concludes that the high-uncertainty scenario is dangerous, and determines the safest action to be a full deceleration for a future left-lane change. The preceding examples show how probabilistic evaluations of predicted worst-plausible evolutions can effectively avoid danger.

VI. CONCLUSIONS

This paper has presented Worst-Case Decision Making, a method that enables the evolution of a dynamic system in an uncertain environment to be predicted. Probability theory is used to model uncertainties, providing a rigorous and intuitive background. The method has been applied to a previously developed deterministic expert system for automated highway driving. The resulting system has been simulated on sample traffic situations involving various degrees of uncertainties. The simulations show that uncertainty can be a major issue when determining an appropriate guidance command. WCDM proves to be an effective tool to correctly analyze the traffic situations and avoid danger. Further simulations will enable the assessment of the performance of the proposed guidance system when applied to multiple vehicles. In particular, these simulations will establish tolerable levels of uncertainty, thus simplifying the choice of an adequate sensor technology.

REFERENCES

[1] *Urban Traffic Systems and Parking*, Transportation Res. Board, Transportation Res. Rec. 1181, Washington, DC, 1988.
[2] A. Reno, "Personal mobility in the United States," in *A Look Ahead Year 2020*, Special Rep. 220, Transportation Res. Board, Nat. Res. Council, Washington, DC, 1988, pp. 369–393.
[3] S. Sheikholeslam and C. Desoer, "Longitudinal control of a platoon of vehicles," in *Proc. 1990 Amer. Contr. Conf.*, San Diego, CA, May 1990, pp. 291–296.
[4] H. Oshizawa and W. Collier, "Description and performance of Navmate, an in-vehicle route guidance system," in *Proc. 1990 Amer. Contr. Conf.*, San Diego, CA, May 1990, pp. 782–787.
[5] P. Tremoulet, "Considerations about the design of an automated-vehicle-collision avoidance system," Human Inform. Processing Group Res. Rep., 89-ES-5, Princeton Univ., Princeton, NJ, Apr. 1989.
[6] R. Stengel and A. Niehaus, "Intelligent guidance for headway and lane control," *Eng. Appl. Artificial Intell.*, Swansea, England: Pineridge Press, vol. 2, pp. 307–314, Dec. 1989.
[7] A. Niehaus and R. Stengel, "An expert system for automated highway driving," *IEEE Contr. Syst. Mag.*, vol. 11, pp. 53–61, Apr. 1991.
[8] ———, "Rule-based guidance for highway driving in the presence of uncertainty," in *Proc. 1991 Amer. Contr. Conf.*, Boston, MA, vol. 3, June 1991, pp. 3119–3124.
[9] "Advanced vehicle/highway systems and urban traffic problems," Science, Education, and Transportation Program, Office of Technology Assessment, U.S. Congress, Staff paper, Washington, DC, 1989.
[10] Siemens AG, *LISB: Leit- und Informationssystem Berlin*, brochure.
[11] H. Okamoto and T. Nakahara, "An overview of AMTICS," in *Convergence 1988—Int. Congr. Transportation Electron: Proc.*, 1988, pp. 219–228.
[12] B. Dobrotin, "Fully automated urban traffic system," Rep. DOT-TST-78-3, pp. A1–A20, 1977.
[13] E. Dickmanns and A. Zapp, "A curvature-based scheme for improving road vehicle guidance by computer vision," *SPIE, Vol. 427, Mobile Robots*, pp. 161–168, 1986.
[14] N. Kehtarnavaz, J. Lee, and N. Griswold, "Vision-based convoy following by recursive filtering," in *Proc. 1990 Amer. Contr. Conf.*, San Diego, CA, May 1990, pp. 268–273.

[15] W. Zhang and R. Parsons, "An intelligent roadway reference system for vehicle lateral guidance/control," in *Proc. 1990 Amer. Contr. Conf.*, San Diego, CA, May 1990, pp. 281–286.

[16] J. Davis, "Adapting radar to the automotive environment," Tech. Rep. Radar Control Systems Corp. (VORAD), May 1987.

[17] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Palo Alto, CA: Morgan Kaufmann, 1988.

[18] P. Szolovits and S. Pauker, "Categorical and probabilistic reasoning in medical diagnosis," *Artificial Intell.*, vol. 11, pp. 115–144, 1978.

[19] D. Stratton and R. Stengel, "Probabilistic reasoning for intelligent wind shear avoidance," presented at the 1990 AIAA Guidance, Navigation & Contr. Conf., Portland, OR, Aug. 1990, paper 90-3437; also *J. Guid., Contr. Dynamics,* to be published.

[20] R. Stengel, *Stochastic Optimal Control: Theory and Application.* New York: Wiley, 1986.

Axel Niehaus received the engineering degree from Ecole Centrale de Paris, France, in 1987, and the M.A. degree in mechanical and aerospace engineering from Princeton University in 1989. He is currently in the revision process of his doctoral dissertation.

Since 1992, he has worked for Matra Automobile, Trappes Cedex, France, as Head of the C.A.E. Department and Project Manager.

Robert F. Stengel (M'77–SM'83–F'93) received the S.B. degree in aeronautics and astronautics from M.I.T. in 1960, and the M.S.E., M.A., and Ph.D. degrees in aerospace and mechanical sciences from Princeton University in 1965, 1966, and 1968, respectively.

He is an Associate Dean of Engineering and Applied Science and a Professor of Mechanical and Aerospace Engineering at Princeton University, where he directs the Topical Program on Robotics and Intelligent Systems and the Laboratory for Control and Automation. Prior to his 1977 Princeton appointment, he was with The Analytic Sciences Corporation, Charles Stark Draper Laboratory, U.S. Air Force, and National Aeronautics and Space Administration. A principal designer of the Project Apollo Lunar Module manual attitude control logic, he also contributed to the design of the Space Shuttle guidance and control system. His current research focuses on system dynamics, control, and machine intelligence. He teaches courses on control and estimation, aircraft dynamics, space flight engineering, and aerospace guidance. He is the author of the book, *Stochastic Optimal Control: Theory and Application* (New York: Wiley, 1986), and has authored or coauthored numerous technical papers and reports.

Dr. Stengel is an Associate Fellow of the AIAA, and a member of the SAE Aerospace Control and Guidance Systems Committee. Professional positions include Associate Editor at Large of the IEEE TRANSACTIONS ON AUTOMATIC CONTROL, North American Editor of the Cambridge University Press Aerospace Series, and member of the Program Council for the New Jersey Space Grant Consortium. He was Vice Chairman of the Congressional Aeronautical Advisory Committee and Chairman of the AACC Awards Committee, and he has served on numerous governmental advisory committees.