Preprints of the 21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Modeling the decision-making in human driver overtaking ⋆

Elis Stefansson * **Frank J. Jiang** * **Ehsan Nekouei** **
Håkan Nilsson *** **Karl Henrik Johansson** *

* *KTH Royal Institute of Technology, Sweden,*
(e-mail: {elisst,frankji,kallej}@kth.se).
** *City University of Hong Kong, China,*
(e-mail: enekouei@cityu.edu.hk)
*** *Uppsala University, Sweden, (e-mail: hakan.nilsson@psyk.uu.se)*

Abstract: We propose models for the decision-making process of human drivers in an overtaking scenario. First, we mathematically formalize the overtaking problem as a decision problem with perceptual uncertainty. Then, we propose and numerically analyze risk-agnostic and risk-aware decision models, which are able to judge whether an overtaking is desirable or not. We show how a driver's decision-making time and confidence level can be primarily characterized through two model parameters, which collectively represent human risk-taking behavior. We detail an experimental testbed for evaluating the decision-making process in the overtaking scenario. Finally, we present some preliminary experimental results from two human drivers.

Keywords: Human decision-making, automated vehicles, overtaking, risk-aware decision-making, driving simulator

1. INTRODUCTION

1.1 Motivation

Transportation networks face the mixed-autonomy challenge wherein automated vehicles, with various degrees of autonomy, are gradually being introduced into road traffic: Lazar et al. (2017); Mehr and Horowitz (2019); Wei et al. (2019). For instance, the first full-sized automated bus will begin to operate in Singapore in conjunction with regular traffic this year, see Wei (2019), and automated minibuses already drive regularly in confined areas in Stockholm: Löfgren (2017). In these mixed traffic environments, there are two types of decision-makers: human drivers and automated vehicles. The safety-level is determined by the interplay between these two types of agents.

The design of decision strategies for automated vehicles requires a thorough understanding of human driver decision-making processes. Here, it is important not just to predict what the human will do, but to *find the underlying mechanisms why the human makes the decision*. Obtaining these key mechanisms is important in its own right, but can also improve decision-making policies for automated vehicles (e.g., enhancing and guiding model-free human prediction models) and to advance driver assistance systems (such as alerting the human driver when an unsafe mode of driving occurs).

Motivated by these observations, this paper considers the decision-making process of human drivers in overtaking scenarios, one of the most dangerous scenarios drivers can encounter, e.g., Hegeman et al. (2004); Wang and Knipling (1994); Barr and Najm (2001). Overtaking scenarios are, from an experimental point of view, rich in the sense that the human is making decisions in a dynamic environment under limited time and subject to risk.

1.2 Contribution

We study the decision-making process of a human driver in an overtaking scenario as illustrated in Fig. 1. Our contributions are three-fold:

(1) We introduce a mathematical formulation of the overtaking problem. In this formulation, an ego vehicle E driven by a human is on a two-way road stuck behind a slow-moving vehicle O, and with an approaching vehicle M in the other direction. E hence has a choice: overtake or wait. We assume E has perfect perception of O, but sequentially receives noisy measurements of M, since M is supposed to be further away.

(2) We propose decision models judging whether E should overtake or not. Each decision model consists of an estimator and a decision rule to overtake or wait based on the current estimate. We consider two decision rules: one risk-agnostic rule that assumes the estimate to be ground-truth, granting an overtaking if the estimate indicates a safe overtaking, and one risk-aware rule that only grants an overtaking if it is confident enough that an overtaking is safe (based on the confidence of the estimate). We analyze the behaviors emerging from these decision models.

(3) We introduce an experimental testbed where a human driver performs overtaking maneuvers in a driving simulator. We explain the setup of the experiment and present some preliminary results.

⋆ This work is supported by the Knut and Alice Wallenberg Foundation, the Swedish Strategic Research Foundation, and the Swedish Research Council.

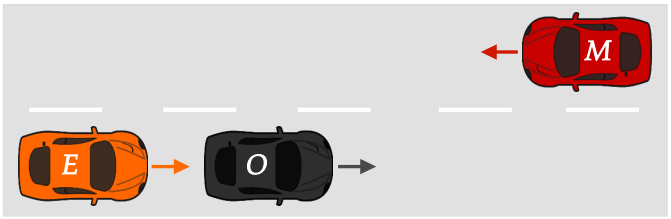

Fig. 1. The considered overtaking scenario.

1.3 Related work

Substantial efforts have been dedicated recently to study and improve the decision-making process of automated vehicles, see Brechtel et al. (2014); Furda and Vlacic (2011); Cunningham et al. (2015). Here, the mixed traffic condition has attracted the use of sophisticated human prediction models for guiding the automated vehicle, e.g., Ziebart et al. (2009); Sadigh et al. (2016); Stefansson et al. (2019). Even though the expressiveness of these models varies, the primary focus is on assisting the automated vehicle by predicting how the human will behave, and not on understanding the underlying mechanisms yielding the human decision. Hence, the assumptions made (e.g., the human has perfect information or knows the automated vehicle's objective) are often left unconfirmed from a human decision-making perspective.

Extensive work has been conducted in the field of psychology to understand the decision-making process of humans. From a psychological perspective, the overtaking scenario involves a choice between a certain undesirable alternative (stay) and a risky alternative (overtake) with both a desirable outcome (successful) and a highly unwanted outcome (unsuccessful). Such choice situations have long been the fruit fly of behavioral decision-making research with numerous systematic behaviors observed, e.g., a general aversion to risky or uncertain alternatives, see Oppenheimer and Kelso (2015), and related models such as Prospect Theory and Decision Field Theory, see Kahneman and Tversky (1979), Busemeyer and Townsend (1993). Previously, human behavior in overtaking scenarios has been studied by statistical analysis of experimental results, see Gray and Regan (2000, 2005); Gordon and Mast (1970). In contrast, we provide a mathematical model for human decision-making in overtaking scenarios based on risk. Finally, the human perception of our model is based on optimal use of the observed information leading to a Kalman filter formulation. Speekenbrink and Shanks (2010) validate human prediction models in dynamic environments based on Bayesian filters, but focus on non-interactive stock market predictions and not on driving scenarios.

1.4 Outline

The paper is organized as follows. The mathematical formulation of the overtaking problem is described in Section 2. Our decision models for the overtaking problem are presented in Section 3. A numerical study of our decision models is performed in Section 4, followed by the description of our experiment and its results in Section 5. Section 6 concludes the paper.

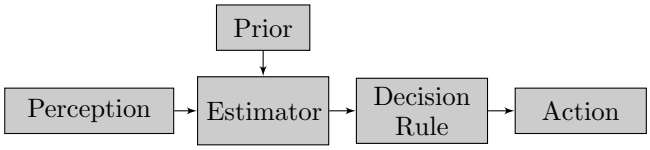

Fig. 2. The underlying structure of the proposed decision models.

2. PROBLEM FORMULATION

Consider an ego vehicle E driving on a two-way road as in Fig. 1, stuck behind a slow-moving vehicle O, and with an approaching vehicle M in the opposite direction. Having a higher desired speed than O, E has a choice: stay behind O or overtake it? This choice should depend on the position and velocity of M, as well as how accurate E estimates these variables. For example, if E's current estimate is inaccurate, it might decide to stay behind O for some time, accumulating more information to obtain a better estimate. However, if E waits too long, M will get too close, prohibiting an overtaking.

Solutions to this overtaking problem are proposed in the next section via decision models judging whether an overtaking is desirable or not. The decision models all make the following standing assumptions:

(i) E has been tailing O for some time with a constant distance and a constant speed v_o (O's speed);
(ii) E has an accurate estimate of the relative position and speed of O;
(iii) E can sequentially obtain noisy measurements of the relative position and velocity of M.

The dynamics are modelled as a discrete-time system indexed by $t \in \mathbb{N}$ (start time $t = 0$) with state $s(t) = [\Delta x(t), \Delta v(t)]^T$, where $\Delta x(t)$ is the longitudinal distance between E and M, and $\Delta v(t) := v_o - v_m$ is the corresponding velocity difference. The state evolves with time step $\Delta t > 0$ according to

$$s(t+1) = Fs(t)$$

where $F = [1, -\Delta t; 0, 1]$. Furthermore, in accordance with Assumption (iii), at each time t, E obtains noisy measurements

$$y(t) = s(t) + \eta(t),$$

where $\eta(t) \sim \mathcal{N}(0, R(t))$ is Gaussian noise with zero mean and covariance $R(t)$.

3. DECISION MODELS

This section considers decision models for the overtaking problem. The considered decision models for the ego vehicle E have the structure in Fig. 2. At each time step t, the estimator, potentially initialized by a prior, is updated with a new (perceptual) measurement of the environment. The updated estimate is then passed to a decision rule which evaluates the decision based on the current estimate. The process ends with an action being taken (overtake or wait).

3.1 Estimator

The estimator assumes that E takes into account the obtained perceptual measurements $\{y(0), \ldots, y(t)\}$, and

uses this information optimally to obtain an estimate of the current state $s(t)$. Mathematically, let $\hat{s}_{t|t'}$ denote the estimate at time t given measurements $\{y(0), \dots, y(t')\}$ ($t' \leq t$), and $P_{t|t'} = \mathbb{E}[(s(t) - \hat{s}_{t|t'})(s(t) - \hat{s}_{t|t'})^T]$ its corresponding covariance error. Our optimality assumption then states that the estimator $\hat{s}_{t|t}$ should minimize the covariance error

$$\min_{\hat{s}_{t|t}} \mathbb{E}[(s(t) - \hat{s}_{t|t})(s(t) - \hat{s}_{t|t})^T]. \tag{1}$$

For a given initial distribution $s(0) \sim \mathcal{N}(\mu, \Sigma)$, the Kalman filter estimate $\hat{s}_{t|t}$ minimizes (1), obtained recursively via the Kalman filter equations with initialization $\hat{s}_{0|-1} = \mu$ and $P_{0|-1} = \Sigma$. In practice, the true distribution $s(0) \sim \mathcal{N}(\mu, \Sigma)$ is unknown to E. The estimator can then instead have an internal guess (μ_0, Σ_0) of (μ, Σ), initializing the Kalman filter with $\hat{s}_{0|-1} = \mu_0$ and $P_{0|-1} = \Sigma_0$, i.e., using the proxy $s(0) \sim \mathcal{N}(\mu_0, \Sigma_0)$. For brevity, we then say that the Kalman filter is initialized with prior (μ_0, Σ_0). Alternatively, the estimator does not have a prior (μ_0, Σ_0), but, for instance, estimates $s(0)$ with the first measurements. One such case we consider is the initialization $\hat{s}_{0|0} = y(0)$ and $P_{0|0} = R$. We then say that the Kalman filter is initialized without prior. In both cases, subsequent estimates are obtained via the Kalman filter equations.

An important aspect with having a prior (μ_0, Σ_0) is that it enables us to capture a priori belief of E, such as bias in the decision model. Since human actions typically depend on past experiences (e.g., the human has learned), the prior serves as a tool to model such features. As an example, an experienced driver might have a very strong prior that the meeting car M will appear right after another meeting vehicle due to previous encounters of similar situations, while an unexperienced driver might exclusively focus on what she sees having a very uncertain prior. In Section 4, we systematically investigate how such priors can affect the overtaking decision.

We next consider our decision rules: the risk-agnostic and the risk-aware decision rules.

3.2 Decision Rules

Risk-agnostic decision rule The risk-agnostic decision rule obtains at each time t an estimate $\hat{T}_m(t)$ of the time $T_m(t)$ when E would meet M during overtaking, and compares it with the time T_o it would take E to overtake O. It then overtakes if $T_o < \hat{T}_m(t)$. More precisely, assuming E needs to cover a relative distance Δx_o when overtaking O, holding a overtaking velocity $v_{to} > v_o$, we get

$$T_o = \frac{\Delta x_o}{v_{to} - v_o}.$$

Here, we have assumed that E can change lanes momentarily and velocity momentarily from v_o to v_{to} and that v_{to} is constant throughout the overtaking, yielding a constant overtaking time T_o. Furthermore,

$$T_m(t) = \frac{\Delta x(t)}{v_{to} - v_m} = \frac{\Delta x(t)}{\Delta v(t) + (v_{to} - v_o)}$$

yields the estimate $\hat{T}_m(t) = \Delta\hat{x}(t)/[\Delta\hat{v}(t) + (v_{to} - v_o)]$. The decision rule is then

$$D(\hat{s}_{t|t}) = \begin{cases} \text{Overtake} & \text{if } T_o < \hat{T}_m(t) \\ \text{Wait} & \text{otherwise.} \end{cases}$$

Risk-aware decision rule The risk-agnostic decision rule overtakes whenever $T_o < \hat{T}_m(t)$, independent of the uncertainty in $\hat{T}_m(t)$. Thus, the decision rule is agnostic to risk and, hence, is potentially unsafe. To account for the uncertainty in the current estimate, the risk-aware decision rule overtakes only when the probability of the event $T_o < T_m(t)$ is high, conditioned on the current estimate $\hat{s}_{t|t}$. More precisely, the decision rule is

$$D(\hat{s}_{t|t}) = \begin{cases} \text{Overtake} & \text{if } \mathbb{P}(T_o + \beta < T_m(t) \mid \hat{s}_{t|t}) > 1 - \delta \\ \text{Wait} & \text{otherwise.} \end{cases} \tag{2}$$

Here, $0 \leq \delta \leq 1$ is the confidence level of the decision, and $\beta \geq 0$ prescribes an additional time margin for the overtaking.

Define the constant row vector and scalar $L := [-1, T_o + \beta]$ and $b := -(T_o + \beta)(v_{to} - v_o)$. The computation process for (2) is then:

(1) Obtain jointly Gaussian distribution of $[Ls(t), \hat{s}_{t|t}]$.
(2) Obtain conditional Gaussian distribution $Ls(t)|\hat{s}_{t|t}$.
(3) Calculate the confidence $\mathbb{P}(T_o + \beta < T_m(t) \mid \hat{s}_{t|t}) = \mathbb{P}(Ls(t) < b \mid \hat{s}_{t|t})$ (using $Ls(t)|\hat{s}_{t|t}$) and overtake if and only if it is greater $1 - \delta$. [1]

Under the assumption that the Kalman filter is initialized with prior (μ, Σ), the conditional Gaussian distribution $Ls(t)|\hat{s}_{t|t}$ is equal to $Ls(t)|y_0, \dots, y_t$. Therefore, the confidence in (2) is equal to $c_t := \mathbb{P}(T_o + \beta < T_m(t) \mid \hat{s}_{t|t}) = \mathbb{P}(T_o + \beta < T_m(t) \mid y_0, \dots, y_t)$, hence optimal, since there is no performance loss working with $\hat{s}_{t|t}$ instead of the measurements. One may also ask if the decision rule is optimal (under certain conditions) given the current confidence c_t. The answer is yes. To see it, abbreviate the events "overtaking successful", "overtaking not successful", and "not overtaking" by OS, ONS, NO, respectively. Let u_E be the utility of E defined by the triple $(u_E(OS), u_E(ONS), u_E(NO))$ and $\mathbb{E}[u_E|a]$ denote expected utility with respect to c_t given action $a \in \{\text{Overtake}, \text{Wait}\}$. That is, $\mathbb{E}[u_E|\text{Overtake}] = c_t u_E(OS) + (1 - c_t)u_E(ONS)$ and $\mathbb{E}[u_E|\text{Wait}] = u_E(NO)$. Consider the decision rule

$$D(\hat{s}_{t|t}) = \begin{cases} \text{Overtake} & \text{if } \mathbb{E}\big[u_E|\text{Overtake}\big] > \mathbb{E}[u_E|\text{Wait}] \\ \text{Wait} & \text{otherwise} \end{cases}. \tag{3}$$

It is easy to see that (2) coincides with (3) provided

$$u_E(ONS) < u_E(NO) < u_E(SO)$$
$$(1 - \delta)u_E(OS) + \delta u_E(ONS) = u_E(NO).$$

Hence, (2) is in this sense optimal being equivalent to the solution of an expected utility maximization.

4. NUMERICAL EVALUATIONS

In this section, we numerically investigate the performance of the decision models. We start with the risk-agnostic

[1] Here and throughout this section, we assume that E and M always travel towards each other: $\mathbb{P}(\Delta v(t) > 0) = 1$.

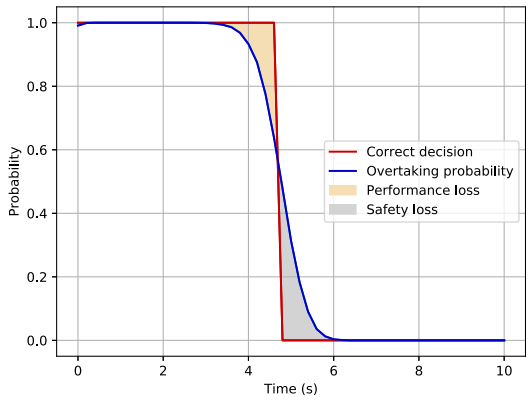

Fig. 3. Probability $\mathbb{P}(T_o < \hat{T}_m(t))$ that the risk-agnostic decision rule with estimator having no prior grants an overtaking as a function of time in blue, and correct decision $\mathbb{1}\{T_o < T_m(t)\}$ in red.

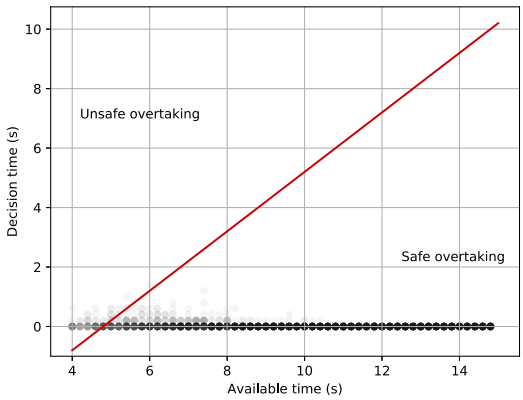

Fig. 4. Scatter plot for the risk-agnostic decision rule with estimator having no prior. The red line marks the border between safe overtakings $(T_o < T_m(t))$ and unsafe overtakings $(T_o > T_m(t))$.

decision rule, showcasing its shortcomings, and then continue with the risk-aware decision rule illustrating how the parameters in the model affect the outcome, in particular having a prior.

4.1 Risk-agnostic decision rule

We combine the risk-agnostic decision rule with an estimator with no prior. The parameters used are found in Table 1, assuming, for simplicity, constant noise covariance $R(t) \equiv R$. To get some intuition, we plot the probability $\mathbb{P}(T_o < \hat{T}_m(t))$ that this decision model grants an overtaking[2] as a function of time (start time zero), seen in Fig. 3, together with the correct decision $\mathbb{1}\{T_o < T_m(t)\}$. If the overtaking probability is lower than the correct decision (i.e., the decision model sometimes does not grant overtaking when it should), then we have a performance loss, while if the overtaking probability is higher than the correct decision (i.e., the decision model sometimes grant overtaking when it should not), then we have a safety loss. Since the rule is risk-agnostic, it should not differentiate between these losses and we hence suspect them to be within the same order of magnitude. Looking at Fig. 3, we see that this is also the case. In particular, the safety loss is a major part of the deviation, which is undesirable.

Table 1.

Δt	$\Delta x(0)$	$\Delta v(0)$	Δx_o	v_{to}	v_o	v_m	R
0.2	550	55	24	30	25	-30	$\text{diag}(80^2, 15^2)$

The safety-loss can also be illustrated in an "available time-decision time"-diagram where the decision time is the first time the decision rule grants an overtaking and the available time is the time when M passes E (if no overtaking). We vary the available time by setting different initial distances between E and M. Running 50 realizations for each time $t \in \{4, 4+\Delta t, 4+2\Delta t, \dots, 15-\Delta t, 15\}$, we obtain a scatter plot seen in Fig. 4, where the red line marks the border between safe overtakings $(T_o < T_m(t))$ and unsafe overtakings $(T_o > T_m(t))$. The risk-agnostic feature is also here present with a large part of undesired, unsafe overtakings for small available times.

4.2 Risk-aware decision rule

We start by investigating the risk-aware decision rule with no prior in the estimator and then continue with the prior case. The setup in Table 1 is used throughout.

Decision rule without prior We start by illustrating the typical behavior of the confidence $\mathbb{P}(T_o + \beta < T_m(t) \mid \hat{s}_{t|t})$ as a function of time t, noting that it fully determines the output of the decision rule (2). One expects the confidence to have a "low-high-low" trend due to few measurements in the beginning (but large amounts of time) and no time in the end (though many measurements). This is also what we generally get, with example in Fig 5a for $\delta = 0.1$ and $\beta = 0$. Here, the black line is the average confidence over 1000 realizations while the grey area denotes one standard deviation, and the correct decision $\mathbb{1}\{T_o < T_m(t)\}$, the confidence threshold $1-\delta$ (see (2)), and the desired decision $\mathbb{1}\{T_o + \beta < T_m(t)\}$ are plotted in red, dashed green, and dashed orange, respectively. We note that varying δ simply corresponds to vertically changing the threshold $1-\delta$ in dashed green. With $\delta = 0.1$, this threshold is low and the rule quickly grants overtakings. The feature is also apparent in a corresponding scatter plot seen in Fig. 5b, with almost immediate decision times and a small portion of unsafe overtakings for small available times. Decreasing δ and thus increasing the threshold $1-\delta$ (increasing the dashed green line in Fig. 5a) results in more risk-averse behavior with longer decision times and less granted overtakings for small available times, as seen in Fig. 5c with $\delta = 10^{-4}$ and $\beta = 0$. In particular, note the increasing trend in decision time for decreased available time, due to the fact that a smaller available time needs more measurements to yield a high confidence.

So far, we set $\beta = 0$. Adding a time margin $\beta > 0$ is equivalent to decreasing the available time by β seconds. In a plot like Fig. 5a, adding $\beta > 0$ shifts $\mathbb{1}\{T_o + \beta < T_m(t)\}$ horizontally to the left by β relative to the correct decision,

[2] Here, "grants an overtaking" means that the decision rule considers it desirable to overtake but does not execute the action.

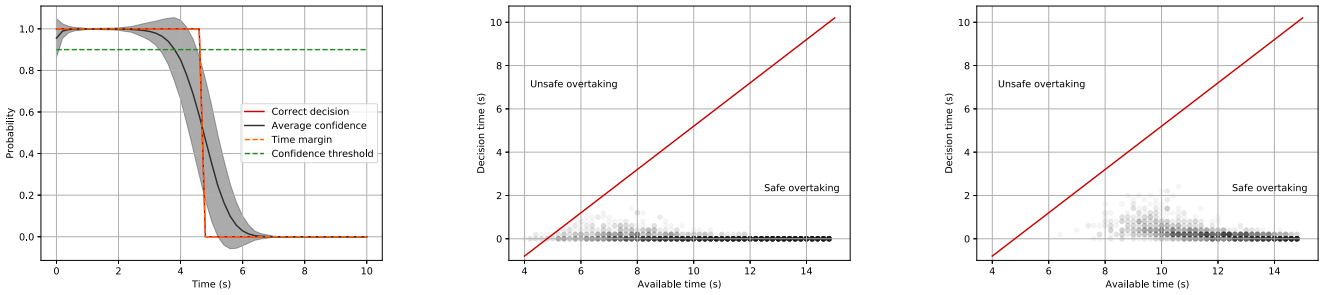

(a) 1000 realizations, $\delta = 0.1$ and $\beta = 0$ (b) $\delta = 0.1$, $\beta = 0$ (c) $\delta = 10^{-4}$, $\beta = 0$

Fig. 5. (a) The average confidence $\mathbb{P}(T_o + \beta < T_m(t) \mid \hat{s}_{t|t})$ in black with one standard deviation in grey for 1000 realizations, with $\delta = 0.1$ and $\beta = 0$. The correct decision $\mathbb{1}\{T_o < T_m(t)\}$, the desired decision $\mathbb{1}\{T_o + \beta < T_m(t)\}$ and $1 - \delta$ are plotted in red, dashed orange and dashed green respectively. (b) Scatter plot as in Fig. 4 for the risk-aware decision rule with no prior in the estimator and $\delta = 0.1$, $\beta = 0$. (c) Scatter plot as in Fig. 4 for the risk-aware decision rule with no prior in the estimator and $\delta = 10^{-4}$, $\beta = 0$.

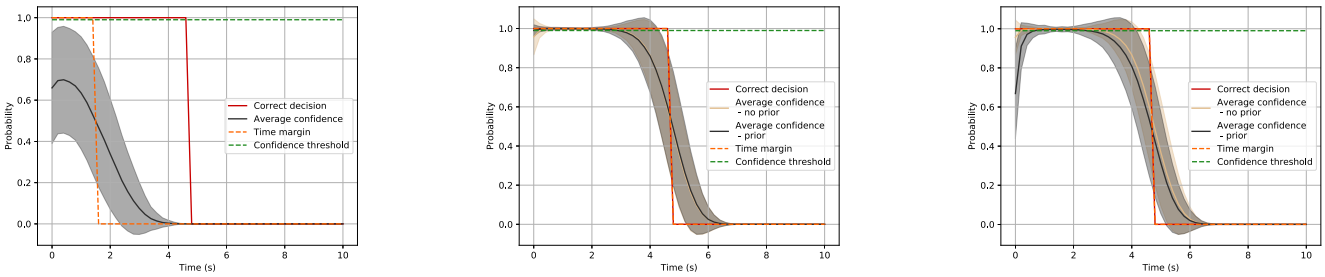

(a) $\delta = 0.01$ and $\beta = 3$ (b) prior $\mu = s(0)$ and Σ_0 (c) prior $\mu_0 = [0, \Delta v(t)]^T$ and Σ_0

Fig. 6. (a) Plot as in Fig. 5a for the risk-aware decision rule with $\delta = 0.01$, $\beta = 3$ and no prior in the estimator. (b) Plot as in Fig. 5a for the risk-aware decision rule with prior $\mu = s(0)$ and Σ_0 as in (4), and $\delta = 0.01$, $\beta = 0$. (c) Plot as in Fig. 5a for the risk-aware decision rule with prior $\mu_0 = [0, \Delta v(t)]^T$ and Σ_0 as in (4), and $\delta = 0.01$, $\beta = 0$.

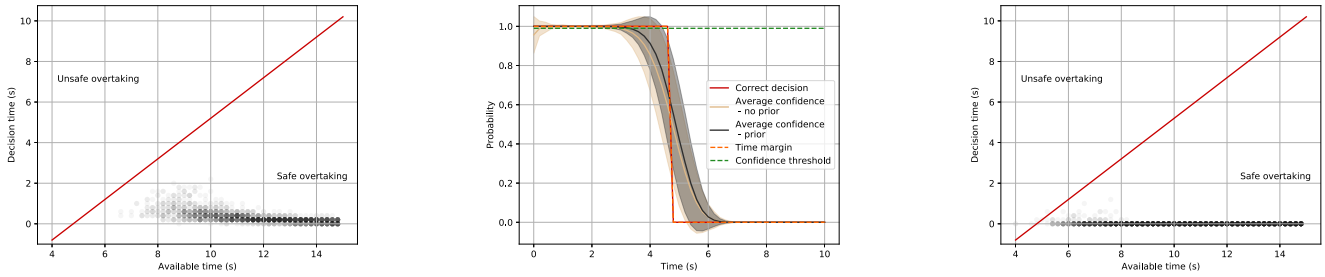

(a) prior $\mu_0 = [0, \Delta v(t)]^T$ and Σ_0 (b) prior $\mu_0 = [2\Delta x(0), \Delta v(t)]^T$ and Σ_0 (c) prior $\mu_0 = [2\Delta x(0), \Delta v(t)]^T$ and Σ_0

Fig. 7. (a) Scatter plot as in Fig. 4 for the risk-aware decision rule with prior $\mu_0 = [0, \Delta v(t)]^T$ and Σ_0 as in (4), and $\delta = 0.01$, $\beta = 0$. (b) Plot as in Fig. 5a for the risk-aware decision rule with prior $\mu_0 = [2\Delta x(0), \Delta v(t)]^T$ and Σ_0 as in (4), and $\delta = 0.01$, $\beta = 0$. (c) Scatter plot as in Fig. 4 for the risk-aware decision rule with prior $\mu_0 = [2\Delta x(0), \Delta v(t)]^T$ and Σ_0 as in (4), and $\delta = 0.01$, $\beta = 0$.

with an example in Fig. 6a for $\delta = 0.01$ and $\beta = 3$. Here, the decision rule has less time to make a decision (having the dashed orange line as the new time limit) with lower confidence as a result. Also, the variation is larger since the first measurements play a more crucial role for estimating if an overtaking is desirable or not. In a scatter plot, the addition of β just shifts the distribution to the right (by β) relative to the case $\beta = 0$.

Decision rule with prior We now consider the risk-aware decision rule and an estimator with a prior (μ_0, Σ_0), fixing $\delta = 0.01$ and $\beta = 0$ throughout. Assuming first a correct mean in the prior, $\mu_0 = s(0)$, a very low Σ_0 results in a confidence following the correct decision $\mathbb{1}\{T_o < T_m(t)\}$ closely (due to an accurate initial estimate of $s(0)$). Higher values of Σ_0 yields more uncertain confidence, though still more certain than having an estimator with no prior, as

Fig. 8. The experimental setup for the overtaking scenario.

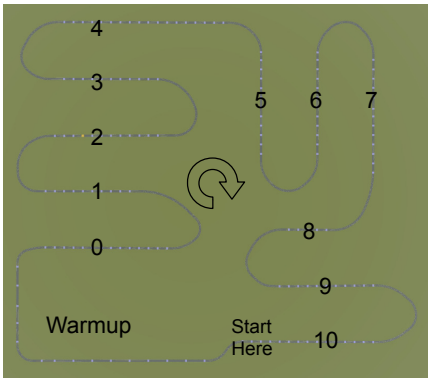

Fig. 9. The experimental road map with an initial warm up region and 11 possible regions to setup an overtaking scenario.

illustrated in Fig. 6b with [3]

$$\Sigma_0 = \begin{bmatrix} 100^2 & 0 \\ 0 & 20^2 \end{bmatrix}. \quad (4)$$

The difference is most apparent in the beginning of Fig. 6b (0-1 s), where the confidence with a prior is higher with lower variation (due to initial guess of $s(0)$) than the no-prior case (having no initial guess). For later times, the two confidences asymptotically approach similar shapes, as the measurements become dominant. In the limit $\Sigma_0 \to \infty$, the two cases become identical.

Consider now the case when we instead have a bias: $\mu_0 \neq s(0)$. Assume first a pessimistic bias expecting M to be closer than the true value. A result with $\mu_0 = [0, \Delta v(t)]^T$ and Σ_0 as in (4) is given in Fig. 6c. Here, the decision model is initially unsure whether an overtaking is desired due to the pessimistic prior, but as measurements accumulate, it approaches the no-prior decision model. The delay in high confidence is also seen in a scatter plot depicted in Fig. 7a.

Finally, we conclude with an optimistic bias $\mu_0 = [2\Delta x(0), \Delta v(t)]^T$ and Σ_0 as in (4). The result is seen in Fig. 7b. The decision model is initially highly certain that an overtaking is desirable, which results in occasional unsafe overtakings seen in Fig. 7c for small available times.

4.3 Summary

The risk-agnostic decision rule is subject to unsafe overtaking behavior due to its risk-agnostic feature. The risk-aware decision rule has safer overtaking with varying behavior depending on the parameters δ, β and potential prior (μ_0, Σ_0): lower δ yields more certain overtakings with longer decision times, especially for small available times (see Fig. 5c); higher β increases the extra marginal time in overtakings corresponding to horizontally shifted distributions in the scatter plots. Finally, the risk-aware decision rule having a prior (μ_0, Σ_0) with pessimistic bias yields longer overtaking times due to initial mismatch while an optimistic bias makes more rush decision with occasional unsafe overtakings for short available times.

5. EXPERIMENTAL EVALUATION

In this section, we outline the experimental setup that we developed and designed for collecting data on human overtaking decision-making processes and present preliminary results from experimental trials on real humans. This experimental setup will be used for an experimental campaign, discussed further in Section 6.

5.1 Experimental Setup

The aim of our experimental setup is to provide a realistic environment in which human drivers end up in scenarios where they need to decide whether to overtake a slow vehicle or not.

For our driving simulator, we built a custom driving environment and API on top of the CARLA simulator (Dosovitskiy et al. (2017)). We interfaced a Logitech G29 steering wheel and pedals with the simulator. We position the human driver's perspective in the simulator relative to the steering wheel such that the driver feels that she is in the cabin of the simulated vehicle. An example of a human driver interacting with the simulator is shown in Fig. 8.

At the beginning of each trial, we first introduce the subject to the simulator and the task we would like her to complete. Then, we allow her to use the simulator for 15 minutes while we record several data points on her interactions.

In Fig. 9, we show the full route of the experiment. The experiment starts with 2 minutes of warm-up where the subject gets used to the controls of the simulator and encounters vehicles in the oncoming direction. Then, the subject begins to enter the different stages of the experiment. In each stage, the subject (E in Section 2) encounters the following entities:

- *a slow vehicle* (O in Section 2) in front of the subject in the same lane, which the subject should try to overtake;
- *a line of blocking vehicles* in the oncoming direction, which should prevent the subject from overtaking the slow vehicle until the subject is tailing the slow vehicle; [4] and
- *a risk vehicle* (M in Section 2) following the line of blocking vehicles at a random distance (serves as our random initial condition in a scenario).

[3] That is, 100 meter standard deviation in $\Delta x(0)$ and 20 m/s standard deviation in $\Delta v(0)$.

[4] These vehicles are not described in Section 2, but are used in the experiment to get vehicle E to tail behind vehicle O; see Assumption (i) in Section 2.

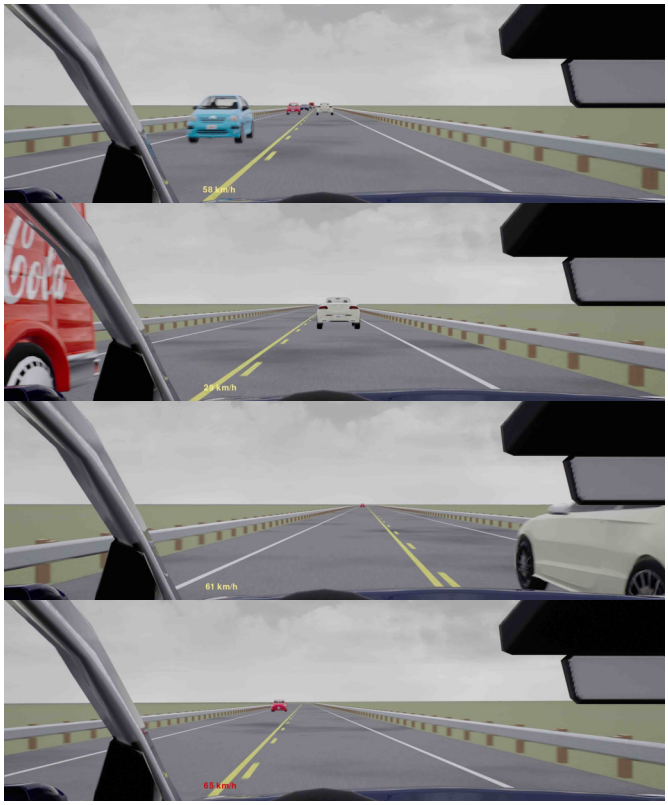

Fig. 10. A series of snapshots from an example overtaking by a human driver. In the first snapshot, we can see the blocking line of vehicles on the left. In the second and third snapshots the slow vehicle is approached and the overtaking maneuver has begun. In the final snapshot, we can see the risk vehicle on the left.

The scenario set up on each stage is illustrated in Fig. 10. There are a total of 11 stages. For each stage, we measure the positions and velocities of each vehicle. Furthermore, we measure the lane crossings of the ego vehicle; these time stamps are used to determine when the human driver made her overtaking decision.

5.2 Preliminary Results

We present two illustrative human driver results. The first human driver is an individual with prior knowledge of the experiment, called Driver A. The second human driver is an individual with no prior knowledge of the experiment, called Driver B. For both drivers, we present their cumulative driving data over all of their interactions with the driving simulator, in Fig. 11. We visualize each of the data sets using the same "available time-decision time"-diagram used in Section 4.

To detect overtakings, we set the start time to be the time when E passes the last vehicle in the line of blocking vehicles. We then wait for E to complete an overtaking maneuver, i.e. first crossing the lane divider into the lane of opposing traffic and then crossing the lane divider again at a position longitudinally greater than O's position. We let the decision time be the time (relative to the start time) when E first crosses the lane divider during the overtaking maneuver. We consider the available time as the time when E meets M longitudinally. Finally, we estimate T_o directly

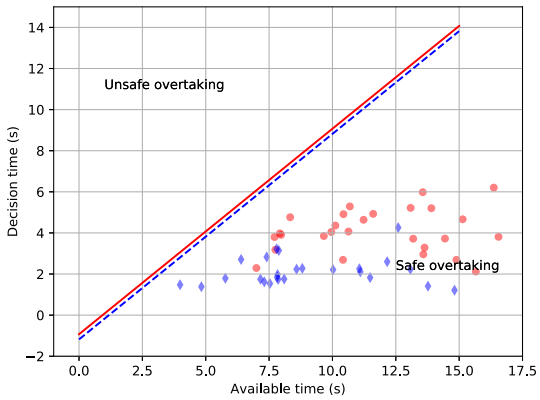

Fig. 11. The "available time-decision time"-diagram for Driver A (red circles) and Driver B (blue diamonds). The line for each subject is determined by the corresponding estimated overtaking time T_o.

by taking the minimum observed overtaking time for the subject.

We collected data on a total of 72 (60) minutes of driving time and 27 (25) overtaking decisions from Driver A (B). We plot the results in red (blue) in Fig. 11, where we observe the following characteristics about the two driver's data sets: [5]

- both Driver A and Driver B never make an unsafe overtaking;
- as available time increases, both Driver A and Driver B take increasingly more varied time to make a decision; however, Driver B exhibits a smaller spread;
- the mean decision time remains fairly constant as available time increases.

We find that the two data sets exhibit phenomenon that are not completely captured by our decision models. In particular, there is more variance in the experimental data at larger available times than in the numerical results. Also, the fairly constant mean decision time indicates that the human decision process may not only involve continuous updates of core variables, but also rely on some all-or-nothing behaviors. However, to determine whether these are real phenomena, we will need to collect a more comprehensive data set. As discussed in the concluding section, we will begin an experimental campaign to obtain this statistically meaningful data set to further investigate the validity of our models.

6. CONCLUSION

This paper considered the decision-making process of human drivers in overtaking scenarios. The overtaking problem was mathematically formalised as a decision problem for the ego vehicle (possibly driven by a human) with perceptual uncertainty concerning an approaching vehicle in the opposite lane. Based on this formulation, we proposed

[5] Note, the available times for each data set vary due to the random initialization of the position of the risk vehicle in each scenario and the driving habits of the individual driver (e.g., some drivers may keep closer to the slow vehicle, some stay further away).

decision models judging whether an overtaking is desirable or not. Each decision model is composed by an estimator and a decision rule, where the latter makes the decision based on an estimate from the former. The estimator is implemented as a Kalman filter, and two decision rules were considered: one risk-agnostic rule that assumes the estimate to be ground-truth, granting an overtaking if the estimate indicates a safe overtaking, and one risk-aware rule that only grants an overtaking when its confident enough that an overtaking is safe (based on the confidence of the estimate). We also presented a comprehensive analysis illustrating how the parameters in the decision models results in different overtaking behavior. Furthermore, we presented a realistic experimental testbed that we use as an evaluation platform for human driving overtaking. Finally, we presented some preliminary results where we begin to see the type of data our experimental platform yields.

Future work includes investigating other estimators (e.g., particle filters), and see how the covariance $R(t)$ can be obtained from perceptual studies in psychology. We will also start an experimental campaign to obtain a statistically significant driving behavior data set and see how the result relates to the current decision models; in particular, we will estimate β and δ in the risk-aware decision rule from the obtained data set and see how well the parameters reproduce the result.

ACKNOWLEDGEMENTS

The authors would like to thank Jerome Busemeyer and John Baras for fruitful discussions about this work.

REFERENCES

Barr, L. and Najm, W. (2001). Crash problem characteristics for the intelligent vehicle initiative. In *TRB 80th Annual Meeting*.

Brechtel, S., Gindele, T., and Dillmann, R. (2014). Probabilistic decision-making under uncertainty for autonomous driving using continuous POMDPs. In *17th International IEEE Conference on Intelligent Transportation Systems*, 392–399.

Busemeyer, J.R. and Townsend, J.T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100(3), 432–459.

Cunningham, A.G., Galceran, E., Eustice, R.M., and Olson, E. (2015). MPDM: Multipolicy decision-making in dynamic, uncertain environments for autonomous driving. In *IEEE International Conference on Robotics and Automation*, 1670–1677.

Dosovitskiy, A., Ros, G., Codevilla, F., Lopez, A., and Koltun, V. (2017). CARLA: An open urban driving simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, 1–16.

Furda, A. and Vlacic, L. (2011). Enabling safe autonomous driving in real-world city traffic using multiple criteria decision making. *IEEE Intelligent Transportation Systems Magazine*, 3(1), 4–17.

Gordon, D.A. and Mast, T.M. (1970). Drivers' judgments in overtaking and passing. *Human Factors*, 12(3), 341–346.

Gray, R. and Regan, D.M. (2000). Risky driving behavior: A consequence of motion adaptation for visually guided motor action. *Journal of Experimental Psychology: Human Perception and Performance*, 26(6), 1721.

Gray, R. and Regan, D.M. (2005). Perceptual processes used by drivers during overtaking in a driving simulator. *Human Factors*, 47(2), 394–417.

Hegeman, G., Hoogendoorn, S., and Brookhuis, K. (2004). Observations overtaking manoeuvres on bi-directional roads. *Advanced OR and AI Methods in Transportation*, 1, 505–510.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.

Lazar, D.A., Coogan, S., and Pedarsani, R. (2017). Capacity modeling and routing for traffic networks with mixed autonomy. In *2017 IEEE 56th Conference on Decision and Control*, 5678–5683.

Löfgren, E. (2017). Stockholm gets Scandinavia's first driverless buses on public road. URL `https://www.thelocal.se/20171227/stockholm-gets-scandinavias-first-driverless-buses-on-public-road`.

Mehr, N. and Horowitz, R. (2019). How will the presence of autonomous vehicles affect the equilibrium state of traffic networks? *IEEE Transactions on Control of Network Systems*.

Oppenheimer, D.M. and Kelso, E. (2015). Information processing as a paradigm for decision making. *Annual Review of Psychology*, 66(1), 277–294.

Sadigh, D., Sastry, S.S., Seshia, S.A., and Dragan, A.D. (2016). Planning for autonomous cars that leverage effects on human actions. In *Proceedings of Robotics: Science and Systems*.

Speekenbrink, M. and Shanks, D.R. (2010). Learning in a changing environment. *Journal of Experimental Psychology: General*, 139(2), 266.

Stefansson, E., Fisac, J., Sadigh, D., Sastry, S., and Johansson, K.H. (2019). Human-robot interaction for truck platooning using hierarchical dynamic games. In *European Control Conference*.

Wang, J.S. and Knipling, R. (1994). Lane change/merge crashes: problem size assessment and statistical description. Technical report, U.S. Department of Transportation.

Wei, Q., Rodriguez, J.A., Pedarsani, R., and Coogan, S. (2019). Ride-sharing networks with mixed autonomy. In *American Control Conference*, 3303–3308.

Wei, T.T. (2019). NTU and Volvo launch world's first full-sized driverless electric bus for trial. URL `https://www.straitstimes.com/singapore/transport/ntu-and-volvo-launch-worlds-first-full-sized-autonomous-electric-bus-for-trial`.

Ziebart, B.D., Ratliff, N., Gallagher, G., Mertz, C., Peterson, K., Bagnell, J.A., Hebert, M., Dey, A.K., and Srinivasa, S. (2009). Planning-based prediction for pedestrians. In *RSJ International Conference on Intelligent Robots and Systems*, 3931–3936. IEEE.