

A model of dyadic merging interactions explains human drivers' behavior from control inputs to decisions

Siebinga, Olger; Zgonnikov, Arkady; Abbink, David A.

DOI

[10.1093/pnasnexus/pgae420](https://doi.org/10.1093/pnasnexus/pgae420)

Publication date

2024

Document Version

Final published version

Published in

PNAS Nexus

Citation (APA)

Siebinga, O., Zgonnikov, A., & Abbink, D. A. (2024). A model of dyadic merging interactions explains human drivers' behavior from control inputs to decisions. *PNAS Nexus*, 3(10), Article pgae420. <https://doi.org/10.1093/pnasnexus/pgae420>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

A model of dyadic merging interactions explains human drivers' behavior from control inputs to decisions

Olger Siebinga ^{a,*}, Arkady Zgonnikov ^a and David A. Abbink^a

^aMechanical Engineering - Cognitive Robotics, TU Delft, Delft, CD 2628, The Netherlands

*To whom correspondence should be addressed: Email: o.siebinga@tudelft.nl

Edited By Lydia Kavradi

Abstract

Safe and socially acceptable interactions with human-driven vehicles are a major challenge in automated driving. A good understanding of the underlying principles of such traffic interactions could help address this challenge. Particularly, accurate driver models could be used to inform automated vehicles in interactions. These interactions entail complex dynamic joint behaviors composed of individual driver contributions in terms of high-level decisions, safety margins, and low-level control inputs. Existing driver models typically focus on one of these aspects, limiting our understanding of the underlying principles of traffic interactions. Here, we present a Communication-Enabled Interaction model based on risk perception, that does not assume humans are rational and explicitly accounts for communication between drivers. Our model can explain and reproduce observed human interactions in a simplified merging scenario on all three levels. Thereby improving our understanding of the underlying mechanisms of human traffic interactions and posing a step towards interaction-aware automated driving.

Keywords: driver modeling, highway interactions, human factors, automated driving

Significance Statement

Interactions between vehicles are an essential part of everyday traffic. However, the mechanisms underlying these interactions are not well understood. Here, we propose a novel model of traffic interactions between two human drivers during highway merges. Our model assumes that communication between drivers and individual risk perception are key mechanisms underlying merging interactions. The model contributes to gaining fundamental knowledge about how humans interact in traffic but could also have practical applications. State-of-the-art automated vehicles use driver models to predict (and “understand”) human behavior; therefore, the model could contribute to developing autonomous vehicles that can better deal with such interactions.

Introduction

Automated driving holds many potential benefits for society (1–3), but, safe and efficient interactions between Automated Vehicles (AVs) and human-driven vehicles remain an open problem (4). Such interactions frequently occur in everyday traffic: at intersections, on roundabouts, and on highways. This work will focus on merging interactions on highways as they are especially intricate due to the high speeds and multiple available options to resolve a conflict (Fig. 1A).

A potential solution to handling such interactions in AVs is through interaction-aware controllers (e.g. (7, 8)). These controllers assume that human drivers unilaterally respond to the AV's behavior and use a model to predict these responses (9). However, real-world merging interactions are inherently *reciprocal* (10): a driver does not only respond to another driver but also influences their behavior through implicit (or even explicit) communication (4, 11). Individual control inputs and decisions of two or more drivers in a merging situation lead to a *joint interaction*

outcome on multiple levels (Fig. 1C): high-level decision-making (negotiating who goes first), acceptable safety margins, and required individual control inputs. This makes real-world merging behavior complex to understand and model, both from an individual and joint perspective (for a real-world example, see Fig. 1B). Interaction-aware AVs should use a model that captures this complexity, which is currently lacking.

Previous work has shown that in merging interactions, *high-level individual decisions* are made to yield or to go (4, 12), which lead to universal joint outcomes in terms of who merges first based on the kinematics of a merging scenario (13). The *safety margins* (e.g. gaps between two vehicles) are the result of joint behavior (14), but at the same time, these gaps are used by individual drivers to communicate their intent (11). *Low-level control inputs* (i.e. acceleration, velocity, and position) are used by individual drivers to communicate (11, 15), thereby playing an essential role in both the *outcome* of the interaction and the *human perception* of other vehicles' behavior. These interrelated aspects

Competing Interest: The authors declare no competing interests.

Received: February 23, 2024. **Accepted:** September 10, 2024

© The Author(s) 2024. Published by Oxford University Press on behalf of National Academy of Sciences. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

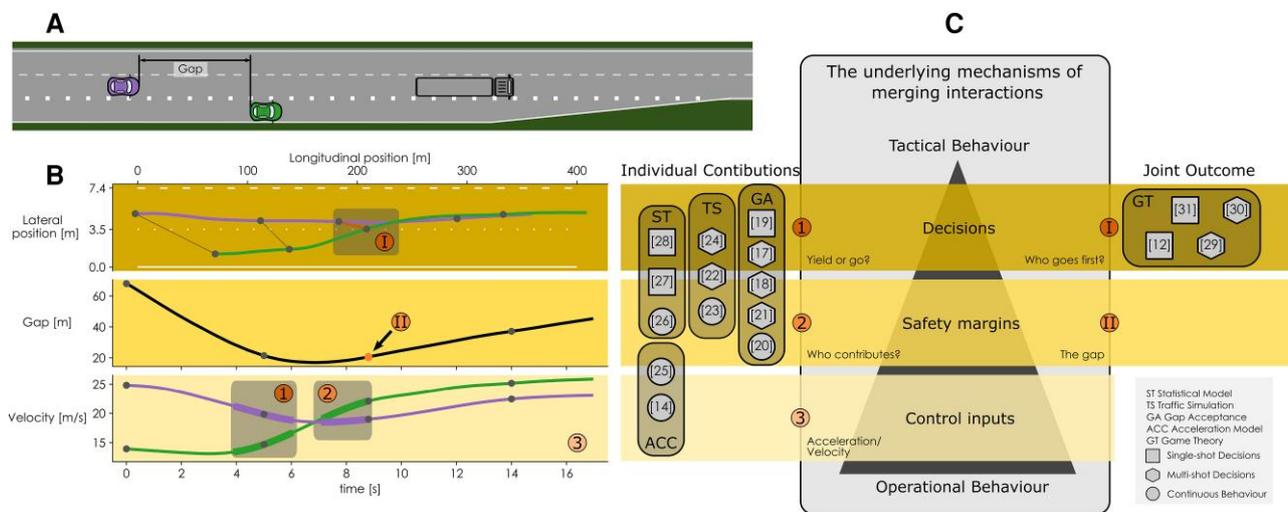


Fig. 1. Models of highway interactions and the aspects of interaction they describe in a merging scenario. A) a typical interactive merging scenario taken from the HighD dataset (5) (dataset 60, vehicles 458 and 468). In this scenario, the driver of the green vehicle wants to merge onto the highway. This vehicle has a position advantage but a significantly lower velocity compared to the purple vehicle driving on the highway. B) the vehicles' position traces, the gap between the vehicles, and the individual velocity traces. In this example, the joint high-level decision is indicated at I: Green merges ahead of Purple (i.e. the green, merging vehicle goes first). Both vehicles individually contribute to this decision by accelerating and decelerating respectively (at 1). After the decision has been made, Green keeps accelerating and thereby individually contributes to maintaining a safety margin while purple stops decelerating (2). The gap between the vehicles, when Green crosses the lane marker (II), denotes the joint safety margin. Finally, the underlying characteristics of the individual vehicle control inputs are depicted here as the total velocity traces (3). We evaluate velocity traces instead of raw accelerations because they are easier to perceive for other drivers and provide more insight into the trend of the driver's actions since they are less noisy. These individual and joint perspectives on the three levels of behavior are also indicated in (C). C) the three levels of behavior in between Michon's *operational* and *tactical* behavior (6). It also shows five modeling strategies for merging interactions, each with examples from the literature. The icons indicate if the models describe a single decision (one-shot), repeated decisions (multishot), or continuous behavior. Every modeling strategy captures part of the overall interactive behavior, but none covers all five aspects. We postulate that a model capturing all three levels of individual and joint behavior simultaneously is likely to have captured the underlying mechanisms of merging behavior.

of individual and joint behaviors are not well understood, and a driver model capturing all aspects is missing.

Our work builds on related work in merging and lane-changing models (merging is often considered a special type of lane change (16)), in which we identified five classes: (i) gap acceptance; (ii) traffic simulation; (iii) statistical; (iv) acceleration; and (v) game theoretic models (Fig. 1C). Gap acceptance models (e.g. (17–21)) describe the decisions made by the individual merging drivers by evaluating available gaps (safety margins) against a personal minimal acceptable gap size. Traffic simulation models often rely on the same gap acceptance theory for making high-level decisions (22–24) and are complemented with acceleration models (e.g. the intelligent driver model (IDM) (14)) to include the control behavior before and after the merging decision. These acceleration models describe individual accelerations and do not include interactive or communicative behavior. We found one acceleration-based model that describes individual drivers' control inputs and safety margins during interactions (25). Statistical models provide a probability that a certain vehicle will merge or change lanes based on naturalistic traffic data. Some include desired safety margins (26, 27), while others do not (28). Finally, game-theoretic models describe the high-level outcome and decision-making of multiple drivers in a single model (e.g. (12, 29–31), see (32) for a review). Game theoretic models assume humans to be rational utility-maximizing agents that do not communicate. However, it is known that these assumptions do not hold for merging drivers (10).

To achieve predictable, legible (33), acceptable, and safe automated behavior, we need to provide interaction-aware AVs with a driver model that covers all three levels: decisions, safety margins, and control inputs. The decision level is important because if automated driving violates the underlying behavioral norms

of human drivers on this level (e.g. it claims the right of way) its behavior will be unacceptable to passengers and other drivers (4). Automated behavior should adhere to acceptable safety margins (on a joint level) and understand how individual drivers keep these margins. This way, AVs will show behavior that is not just safe but is also perceived as safe. Finally, understanding the subtleties of gaps, positions, and velocities can help AVs understand the communication from other drivers and act accordingly. Since merging is a reciprocal interaction, such a model should capture the joint behavior, not just that of a single driver responding to their environment (10).

Besides direct applications in interaction-aware AVs, a complete model of merging interactions could also prove to be a valuable step towards theories and a better fundamental understanding of human interactive capacities and behavior in general (34). A joint driver model could be used to understand and investigate how drivers perceive the behavior of others, how they communicate, and how they negotiate a safe solution in general traffic interactions. This fundamental understanding is needed to design automation that can interact in a natural manner. In a review of automated interactive traffic behavior, Brown et al. (4) stated: "Designing systems that can understand and react to such [implicit traffic] communication will rely upon developing an understanding of that communication beyond statistical regularity." Black-box trajectory prediction models (e.g. (35–37)) do not provide insight into the underlying mechanisms of merging interactions and will thus not suffice for this purpose. If a model of merging interactions succeeds in capturing the underlying mechanisms of merging interactions, which is likely if it captures behavior across multiple levels, it could generalize to other interactive traffic scenarios, helping us gain insight into the fundamentals of interactive human driving behavior.

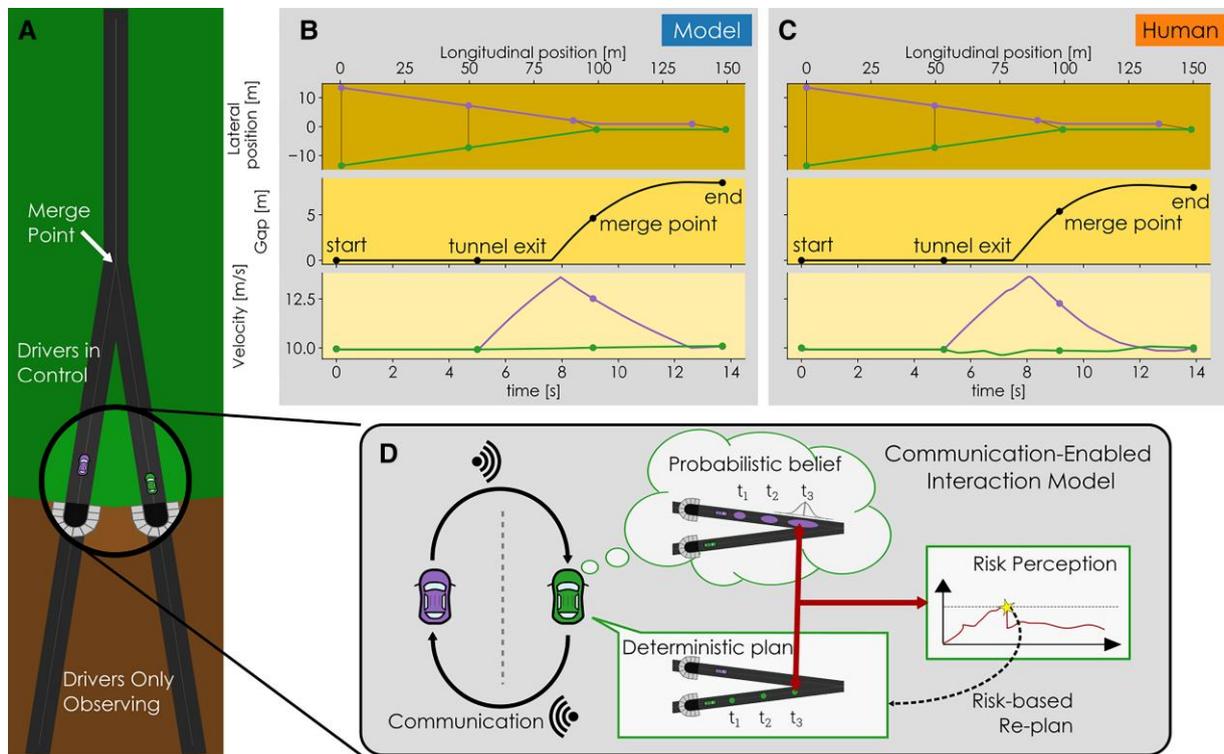


Fig. 2. A) the simplified merging scenario used both in model simulations and the experiment in a coupled top-down-view driving simulator (13). Two vehicles start on different roads, which they follow to a single merge point of those roads. They start in a tunnel where the driver (i.e. either human participants or the model) can observe both vehicles and their initial velocities but have no control yet. After exiting the tunnel, the drivers can now control the vehicles' acceleration (no steering control is needed or available). Beyond the merge point is a short road where the vehicles follow each other. B, C) typical examples of human and model interactions in this scenario (participant pair 3, the model behavior resulted from a fit on the human behavior across all conditions and all trials; the model trial shown here was not fitted specifically to this human trial). D) The CEI framework. For clarity, the panel only visualizes the three model components for the green vehicle, but the model is symmetrical, so the purple vehicle has the same three components. Each driver has a *deterministic plan* for their own behavior and a *probabilistic belief* of the positions of the other driver in the near future. Combined, the plan and belief result in a continuous perception of risk. If this perceived risk exceeds a *risk threshold* the driver alters their plan to return the risk under the threshold. Each driver *communicates* their plan (intention) implicitly (e.g. through vehicle motion) to the other driver, who bases their belief on the received communication. Thus, this communication links one driver's plan to the belief of the other driver.

The main contribution of this work is a novel computational model for a simplified merging scenario with human drivers, based on the Communication-Enabled Interaction (CEI) framework (10). The model assumes that drivers have a deterministic plan for the near future and form a probabilistic belief about another driver's intentions based on implicit communication. The plan and belief result in a perception of risk. If this risk exceeds a personal threshold, a driver updates their plan to get the risk under control. We validate our model on empirical data collected in a top-down view driving simulator with pairs of drivers (13, 38). Our model accurately describes the (qualitative and quantitative) control input characteristics, safety margins and high-level decisions (i.e. who goes first?) of human drivers. It captures differences in individual contributions of drivers and joint behavior. Finally, our model does not assume human rationality and explicitly incorporates communication between drivers as one of the fundamental aspects of interactions, making it the first merging interaction model to avoid these common game-theoretic assumptions.

Results

Simplified merging scenario and experiment

To validate our model of merging interactions between two human drivers, we used data previously gathered in an experiment using a simplified merging scenario (Fig. 2A) in a coupled, top-down view driving simulator (13, 38). The scenario simplifies

merging by simulating two roads (or lanes) that merge into one at a single point. The vehicles start in a tunnel where the drivers can only observe both vehicles traveling at constant velocity to facilitate velocity perception before the interaction. Once both vehicles have exited the tunnel, the drivers gain control over the accelerations of their vehicles (steering is not possible) to resolve a merging conflict. The drivers (nine pairs of participants) were instructed to maintain their initial velocity yet prevent a collision. The experiment used 11 experimental conditions that would end in a collision if the vehicles kept their initial velocity (10 repetitions per condition). A condition consisted of a combination of initial relative velocity ($-0.8, 0.0, \text{ or } 0.8 \text{ m/s}$) and the projected headway at the merge point if both vehicles would continue their initial velocity ($-4, -2, 0, 2, \text{ or } 4 \text{ m}$). The names of the conditions denote *projected headway_relative velocity* (e.g. 4_-8). Positive numbers indicate an advantage for the left driver. We refer to an individual driver as either the left or right driver throughout all trials, based on their physical location during the experiment. The scenario was completely symmetrical (i.e. there was no right of way). Note that from the drivers' view, they randomly perceived approaching the merge point from the left or right side of the track—to account for potential biases due to traffic rules. They were seated in the same room but could not see each other or communicate in any other way than via vehicle motion. For a more extensive description of the experimental protocol, see (13, 38). Figure 2C shows a typical trial outcome with human participants.

Model and CEI framework

We created a novel model based on the CEI framework (10) to describe the joint behavior of a pair of merging drivers. At the core of the CEI framework lies the idea that drivers communicate their plans (intentions) to others using implicit or explicit communication. Empirical evidence has shown that this kind of communication plays an important role in traffic interactions (e.g. (11, 39)), which to date has been absent in interactive driving models. We assume that drivers form a probabilistic belief about the other driver's future movements (intent) based on this communication. Combined with their own deterministic plan, this belief underlies a driver's perceived risk. We assume that if the risk exceeds the driver's individual risk threshold, they will unilaterally alter their plan to get the risk under control. The CEI framework (10) describes this overall structure consisting of four modules: communication, plan, belief, and risk perception. The model proposed here instantiates the CEI framework by implementing these four modules.

The methods section provides a full specification of all modules of our model (plan, communication, belief, and risk), yet we give a summary here to help interpret the results. In our model (Fig. 2D), drivers plan a deterministic trajectory (i.e. a set of waypoints) by optimizing comfort and speed over a time horizon. They communicate this plan implicitly through vehicle kinematics (current position, velocity, and acceleration). Drivers' velocity perception is assumed to be noisy. The drivers' belief about the actions of the other vehicle is represented as a set of probability distributions for the other vehicle's positions at specific times in the future. The recent behavior of the other vehicle influences the variability in the belief (i.e. inconsistent behavior increases the variance). The perceived risk is calculated by evaluating the probability that the positions of the ego and the other vehicle (plan and belief) overlap (i.e. the probability of a collision).

We assume every driver has two dynamic risk thresholds: an upper threshold and a lower threshold. The plan is updated either when (i) the upper threshold is exceeded (to prevent a collision) or when (ii) the perceived risk stays below the lower threshold for a certain amount of time τ (to revert to "normal" behavior when the conflict is resolved). Both thresholds are dynamically adjusted by an incentive function reflecting traffic rules and customs. The rationale is that two drivers perceive the same amount of risk, but traffic rules and customs provide a higher incentive for one of them to act. For example, the following vehicle is usually responsible for preventing collisions in a car-following scenario.

Model simulations

The model uses 10 parameters designed to reflect the scenario which were equal for all simulations. We used a grid search to find individual risk threshold parameters to describe the nine pairs (18 drivers) from the experiment (13) (i.e. one upper and lower threshold per driver is used across conditions). The incentive functions are the same for all drivers and were fitted to all experimental data using linear regression. We simulated the same number of trials as in the experiment: 9 participant pairs, 11 conditions, and 10 repetitions of each condition per pair (990 total trials). The model simulations run faster than real-time with an average run time of 2.6 s (Intel Xeon E5 quad-core) for an average real-time duration of 14.2 s. An example of a simulated trial for participant pair 3 can be found in Fig. 2B. All simulated and experimental data can be found online (40, 41), as well as a script to playback individual trials in an interactive GUI (42). For quick access, seven video captures of such playbacks can be found on the GitHub page with the repository (43).

In the remainder of this section, we will evaluate the model behavior on the three behavioral levels presented in Fig. 1C, both for individual and joint behavior. Some trials ended in a collision; these are excluded from the results because they represent edge cases. Collisions happened infrequently and in all conditions for human drivers (28/990) and model simulations (29/990). The [Supplementary material](#) contains more details on how the collisions were distributed over conditions and a qualitative analysis of the phenomena that lead to collisions.

Characteristics and magnitude of control inputs

Empirical evidence showed that human drivers use intermittent piece-wise constant acceleration control to solve merging conflicts in the simplified merging scenario (13). This type of control results in piece-wise linear velocity patterns (roughly triangular in the plots) that indicate clear decision moments when drivers change their control input to help resolve the merging conflict. The model's control behavior is qualitatively similar to the human driver's (Fig. 3A); it replicates the characteristic patterns in the velocity plots.

Pair 3 is highlighted in Fig. 3A to facilitate comparison. In this pair, the left human driver mostly decelerated at the tunnel exit to prevent a collision, although in some cases, they accelerated. The model replicates this behavior. In one trial (each), the left model and left human maintained their initial velocity for about one second before accelerating. Both the timing of the replanning (i.e. the location of the peak of the velocity profile) and the absolute maximum deviation from the initial velocity were consistent between the left driver and the left model. Contrary to the left driver, the right driver in pair three barely acted to mitigate the risk. Only in one case (each), the right model and the human driver decelerated at the tunnel exit to prevent a collision. This happened in the same trial where the left driver delayed their initial response and then accelerated for both the model and the human driver pair.

The magnitude of control inputs applied by each driver was consistent between the model and human behavior for most pairs (Fig. 3B). Some drivers consistently provided very little input (e.g. the left driver in pair 2), while others used higher input levels (e.g. the right driver in pair 7). The model accurately reflected this quantitative difference through the personalized risk thresholds. However, in some specific conditions for specific pairs, the model produced different average behavior caused by outliers (e.g. pair 1, left driver, condition 0_8). In this specific example, the simulated left driver came to a complete stop in 2 out of 10 trials. In both cases, the simulated left driver initially accelerated but quickly changed strategies and started to brake. At this point, the right driver had already responded to the initial acceleration and had also started braking. Thus, both vehicles were braking, and the only safe solution was for the left driver to come to a complete standstill. This sequence of events can be understood as a miscommunication where the left driver interpreted the right driver's behavior as a reason to change strategies, while at the same time, the right driver shows a (delayed) response to the initial strategy of the left driver. These miscommunications also happened with human drivers (e.g. in Fig. 3A, the right human driver sometimes decelerates briefly before accelerating and going first). However, with human drivers, these trials ended in a collision (see the [Supplementary material](#) for details) or with a less extreme maximum deviation; complete standstills do not occur in the human data. This difference could be due to the velocity perception noise or how communication is translated to a belief in

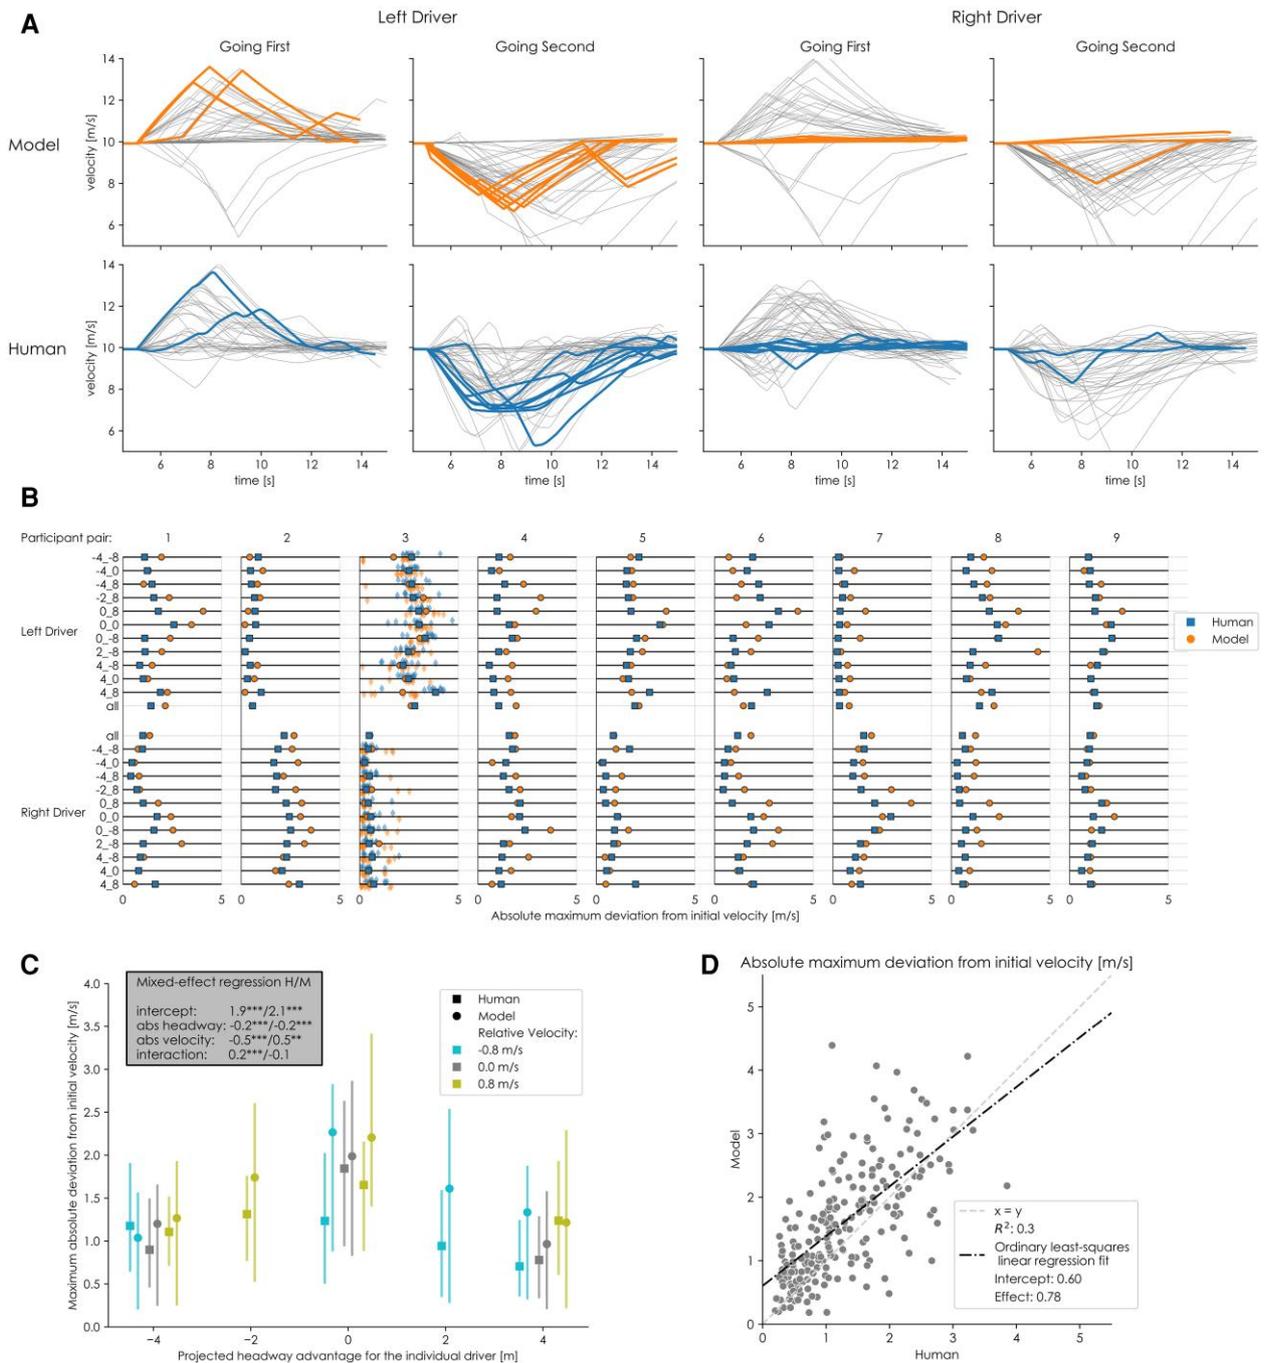


Fig. 3. An overview of a comparison of the control inputs performed by nine pairs of human drivers, and by the model fit to capture these pair’s merging behavior. Trials that ended in a collision were excluded here. A) All velocity traces for the model (top row) and human drivers (bottom row) in condition 0_0 (i.e. no initial velocity or position difference between the vehicles), for each driver (left/right) in the pair, divided into the high-level outcome of the trial (i.e. which driver merged first). To facilitate the comparison for a single pair, the velocities of pair 3 are highlighted. B) mean absolute maximum deviation from the initial velocity for the left and right driver in all participant pairs for all experimental conditions. We use the absolute deviation to compare conditions because a single condition can contain both outcomes within one pair (left going first and right going first, e.g. see A). Thus a single driver could accelerate in some trials and decelerate in others within one condition. The minimum and maximum deviations will be studied separately in section 2 as individual contributions to the decision. For pair 3, all underlying data are visualized (10 trials per condition). One outlier (9.5 m/s, model trial, condition 4_0) is not shown. C) mean absolute maximum deviation from the initial velocity, aggregated over all drivers for different relative velocities and projected headways. The error bars represent interquartile ranges. The headway and velocity values in (C) are shown from the perspective of the individual driver. This means that a projected headway of 4 m represents an advantage from the driver’s perspective, while -4 m represents a disadvantage. The inset indicates coefficients of mixed-effects linear regression models predicting the mean absolute maximum deviation from initial velocity as a function of projected headway and relative velocity (*: $P \leq 0.05$, **: $P \leq 0.01$, ***: $P \leq 0.001$). Full results of the statistical analyses are available in the [Supplementary material](#). D) the relationship between human and model behavior for all participant pairs and all conditions (i.e. all points from B).

the model. However, since these are edge cases that happen infrequently, a more detailed investigation is needed to understand these miscommunications fully.

Across all participants, both human drivers and the model used lower-magnitude inputs with increasing absolute projected headways (Fig. 3C). Absolute relative velocities, however, had opposite

Table 1. Personal base values for the upper and lower thresholds for each driver.

Pair	Driver	θ_l	θ_u
1	Left	0.165	0.495
	Right	0.260	0.562
2	Left	0.245	0.635
	Right	0.058	0.493
3	Left	0.058	0.488
	Right	0.245	0.631
4	Left	0.183	0.537
	Right	0.201	0.524
5	Left	0.113	0.498
	Right	0.269	0.585
6	Left	0.246	0.550
	Right	0.161	0.546
7	Left	0.320	0.736
	Right	0.201	0.522
8	Left	0.165	0.525
	Right	0.246	0.586
9	Left	0.178	0.519
	Right	0.227	0.543

effects on human and model behavior. This can partly be explained by the fact that only for human behavior there is a significant interaction effect. We believe the absence of this effect in the model stems from the independent effects velocity and position have on the belief (i.e. there is no interaction effect in the belief construction). The origins and inner workings of this phenomenon in human behavior are unknown. In general, there is a strong correlation between the model's and human drivers' inputs (Fig. 3D) across all conditions and participants. The model exhibited slightly higher maximum absolute deviations from the initial velocity than human drivers (difference 0.31 m/s, 95% CI [0.23 – –0.40 m/s]). We believe this difference is due to the earlier explained outliers.

Safety margin in terms of the size of the gap between the vehicles

Safety margins can be evaluated individually or at a joint level (Fig. 1); drivers individually contribute to achieving a certain realized safety margin (i.e. gap) on the joint level. These individual contributions can be observed from the absolute control input behavior shown in Fig. 3B. We use the absolute deviation from the initial velocity to compare conditions because in some conditions drivers accelerate in some trials and decelerate in other trials.

In participant pair 3, the left driver mostly contributed to the safe solution; the right driver did not greatly deviate from their initial velocity (Fig. 3B). The model replicates these unequal contributions for this and other pairs (e.g. pairs 2 and 7, and to a lesser extent in pair 5). For the other participant pairs, keeping the safety margin is more of a joint effort, a phenomenon described by the model as well. The model reflects these differences in individual contribution through the baseline risk threshold levels (Table 1). Drivers with a lower tolerance for risk (i.e. upper risk threshold) will act to mitigate their perceived risk, while drivers with a higher tolerance will remain passive. Drivers in pairs with equal contributions (e.g. pairs 4, 6, and 9) have similar upper risk thresholds, while drivers in pairs with unequal contributions (e.g. pairs 2, 3, and 7) have larger individual differences (Table 1).

In some human pairs, the drivers contribute equally to the safety margin, but their relative contributions differ for different conditions (Fig. 3B). For example, in pair 6, both drivers always act. However, in the conditions with a projected headway advantage for the left driver (i.e. a positive number), the right driver

tends to do more, while for the negative projected headways, the left driver acts. This aligns with the assumption that the following driver has more incentive to resolve a conflict. The model captures this behavior (Fig. 3B) through the incentive functions. However, in pair 6, the differences between conditions are smaller in the model than in the human behavior. The limited extent to which the model shows this phenomenon can be attributed to the assumption of identical incentive functions for all pairs. Because only some human pairs show this type of behavior, the averaged incentive functions reduce the extent to which this phenomenon is visible in the model. A method to estimate individual incentive functions could improve this in future versions of the model.

In terms of joint gap-keeping behavior, the human data showed no substantial qualitative differences between pairs (Fig. 4A); the effects of kinematic conditions on the gap are significant but small (Fig. 4B). There is much variability within pairs, where gaps range between 0–9 m within one condition (pair 3, Fig. 4A) and within conditions, with interquartile ranges of 2–3 m (Fig. 4B). This large variability results in a limited correlation between the model and human behavior, where the model overestimates some smaller gaps and underestimates some larger gaps (Fig. 4C). However, overall the model keeps gaps that are comparable in size to human behavior for all participant pairs and all conditions (Fig. 4B). The mean gap for the model was only slightly larger than that of human drivers (4.8 m vs. 4.5 m; difference 0.3 m, 95% CI [0.09–0.51 m]).

Decisions of who merges first and who contributes

Finally, there is the high-level outcome of a merging interaction, which on a joint level can be summarized by the answer to the question: "Who merged first?" In human merging interactions, the probability that a driver merges first increases with their projected headway advantage and decreases with their relative velocity advantage (13). The model replicates these effects (Fig. 5B), although the velocity effect is smaller for the model than for human drivers. The difference in this effect size is especially evident in conditions with pure velocity differences (i.e. a 0 m projected headway) (Fig. 5B). In these conditions, the slower vehicle (i.e. the one that merges first most often) approaches the merge point ahead of the other vehicle (so that they arrive simultaneously). Potential explanations for the discrepancy in effect size between the humans and the model could be that humans systematically underestimate velocity differences or that following vehicles prefer braking over accelerating to prevent collisions. The noise on velocity and the cost of accelerating or decelerating are both assumed to be symmetrical in the model. The similar effect sizes indicate that our proposed combination of kinematics-based probabilistic beliefs and the concept of risk-based control in individual drivers is a strong potential explanation of the underlying principles that govern the high-level outcome in merging interactions.

The intercept of the logistic regression (Fig. 5B) quantifies the asymmetry between outcomes (and participants). The model's intercept is closer to zero than the human drivers', although this effect is insignificant (Fig. 5B). This can be explained by the fact that the model uses the same symmetrical functions for the belief and incentive across all drivers; therefore, the high-level outcome can be expected to be symmetrical as well (i.e. have an intercept of 0.0).

The individual contributions of the drivers to the high-level outcomes can be seen as the question for an individual driver to "go or yield"; i.e. the decision to accelerate or to brake (Fig. 5A). As with the individual contributions to the safety margins, some

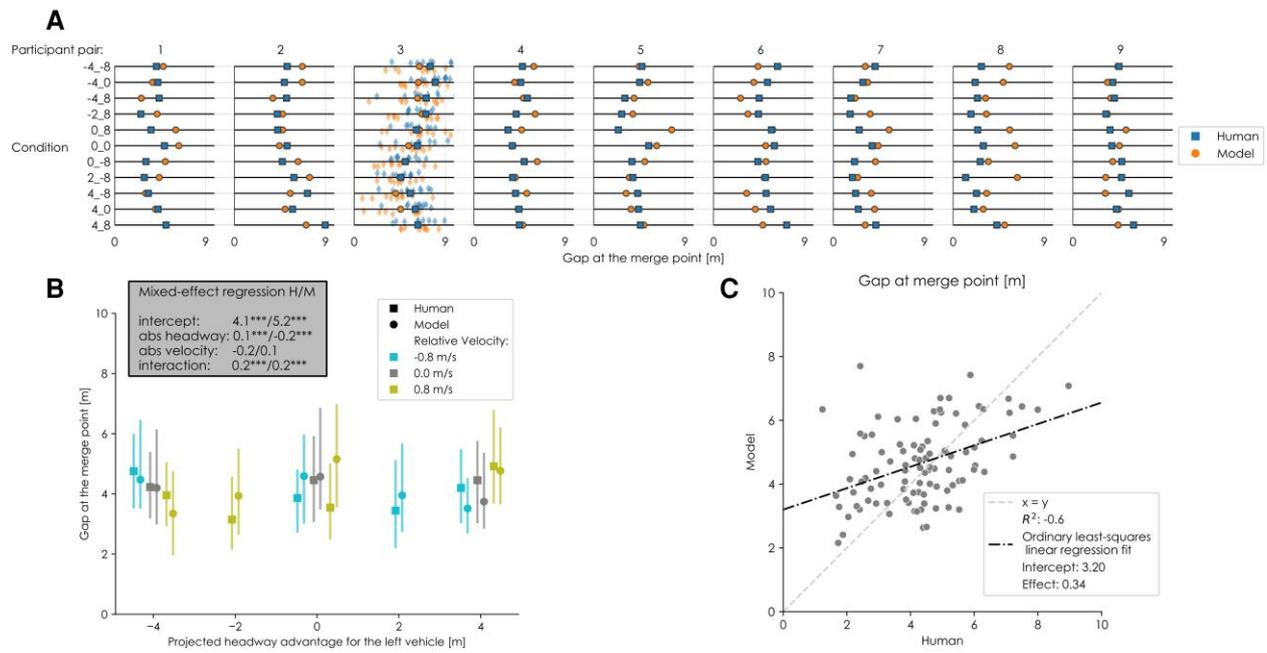


Fig. 4. An overview of the joint gap-keeping behavior of the model and human drivers. A) Mean gap (safety margin) at the merge point for all participant pairs and all conditions, excluding collisions. For pair 3, all data points (i.e. all trials) are shown. B) Mean gaps per condition (aggregated over all drivers); the error bars represent interquartile ranges. The inset indicates coefficients of mixed-effects linear regression models predicting the mean gap as a function of projected headway and relative velocity (*: $P \leq 0.05$, **: $P \leq 0.01$, ***: $P \leq 0.001$); full results of the statistical analysis are available in the [Supplementary material](#). C) The relationship between mean gaps of human drivers and those produced by the model for all participant pairs and all conditions (i.e. all points from A).

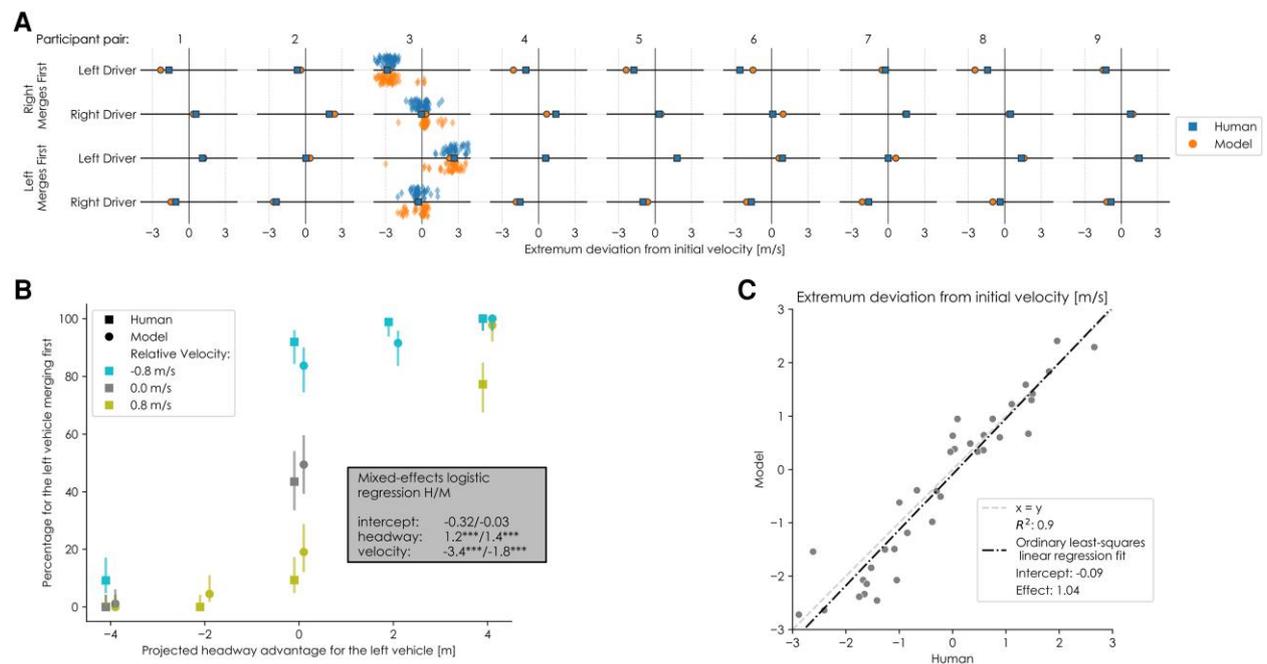


Fig. 5. An overview of the decision-making behavior of the model and human drivers. A) Mean individual contributions of drivers to the high-level outcome of a trial as the maximum or minimum deviation from the initial velocity (i.e. amount of acceleration or deceleration); collisions are excluded. All trials are shown for a representative participant pair (pair 3). B) Joint interaction outcome (i.e. who merged first) for all pairs in every condition. The error bars represent the 95% confidence intervals. The inset indicates coefficients of mixed-effects linear regression models predicting the mean gap as a function of projected headway and relative velocity (*: $P \leq 0.05$, **: $P \leq 0.01$, ***: $P \leq 0.001$); full results of the statistical analysis are available in the [Supplementary material](#). C) The relationship between human and model behavior for all participant pairs and all conditions (i.e. all points from A).

drivers consistently contribute very little to the high-level outcome (e.g. the left driver in pair 2). The model reflects this phenomenon for multiple drivers using their individual thresholds

(left in pairs 2 and 7, right in pairs 3, 5, and 8). However, some drivers only contribute to the high-level outcome when they go second (e.g. the right driver in pair 1), which is reflected by the model

through the incentive function. Finally, some drivers always contribute to the outcome (mostly in interactions with drivers that do nothing, e.g. the left driver in pair 3). The model reflects all these three qualitative phenomena (Fig. 5A). Quantitatively, there is a strong correlation between the humans' and models' decision behavior (Fig. 5C).

Besides the individual decisions that lead to a joint high-level outcome, the model also describes how long the drivers take to reach a decision on a safe outcome. This duration can be measured with the Conflict Resolution Time (CRT) (38). We found that the model captured the previously observed relationship between initial kinematics and CRT (see [Supplementary material](#) for details).

Discussion

We presented a model based on the CEI framework (10) that accurately describes driver behavior in a simplified interactive merging scenario (Fig. 2A). Our model captures the driver behavior on three levels (Fig. 1C): the control input behavior of individual drivers, the safety margins kept by pairs of drivers and how individual contributions establish these, and the high-level decisions of individual drivers (i.e. to merge or yield) and the pair (i.e. who goes first?). Because the model quantitatively and qualitatively captures individual and joint driver behavior on all three levels we consider it likely that the underlying mechanisms of the model (communication-based belief and risk-based replanning) correspond to the mechanisms underlying human interactive driving behavior.

These underlying mechanisms bear a resemblance to mechanisms previously used in models of traffic interactions. The communication-based belief is related to the concept of Theory of Mind (ToM) (44), used in other models of traffic interactions (e.g. (45, 46)). The main difference is that ToM assumes humans to have an internal representation of the motivations of others, while our model uses a more basic belief of future kinematics; i.e. our model does not care if another driver prioritizes speed over safety, it observes a higher velocity and updates a kinematic belief.

Risk-based replanning is a mechanism previously used by Kolekar et al. (47) to model isolated driver behavior in seven real-world driving scenarios. Our model extends the concept of re-planning when the perceived risk exceeds a threshold to an interactive scenario. However, our definition of risk (perceived probability of a collision) is much more simplified than the Driver Risk Field (DRF) used by Kolekar et al. The DRF considers the risk posed by different events (e.g. going off-road or colliding with a tree) and includes steering behavior. Our simplified scenario does not require such a sophisticated definition. However, including the DRF in a new model based on the CEI framework could enable the modeling of real-world merging scenarios.

Finally, the prevalent approach to modeling multiple drivers in interactions is game theory (e.g. (46, 48–51)). Our results have shown that including multiple drivers in a single model (rather than modeling a single driver that responds to their environment) is important because the same joint behavior can stem from multiple individual contributions, and both drivers continuously update their behavior based on the other driver's actions. However, for traffic interactions between vehicles and pedestrians, it has been shown that using game theory to optimize a short-term payoff value is not enough to explain the complex phenomena observed in the real world; instead, a range of more complex mechanisms such as a ToM and implicit communication were needed (45).

Our model and the CEI framework can have important further implications on multiple fronts. The model could be used to further improve our understanding of interactive driving behavior

for the development of automated driving technologies, and our model could potentially be generalized to other scenarios with traffic interactions.

The underlying mechanisms of the model enabled it to replicate human driving behavior on multiple levels. Therefore, the model could help researchers to understand better how these mechanisms function in human behavior (34). For example, implicit communication (through vehicle movements) and how it influences a driver's belief is observed often but only partly understood (11, 39). The same holds for the perception of risk, which has been investigated for isolated drivers (47, 52), but our model could help extend this to interactions. Our model could facilitate research in these directions, leading to an increase in fundamental knowledge.

This knowledge could facilitate more practical applications, such as the design of movements that convey a clear message about the intent of an automated vehicle (4). These movements could consider other drivers' expectations regarding high-level outcomes and communicative actions. Matching the automated vehicle behavior with expectations might increase behavioral acceptance, although this should be investigated further. Second, the model might be used to inform the real-time interaction planning of automated vehicles. In particular, it could inform the AV about the potential future actions of other drivers (as prediction models are generally being used (7, 8, 53)). Finally, our model could also be valuable in the development phase of automated behavior by being part of an interactive and dynamic environment for benchmark testing where models are used to evaluate the behavior of autonomous vehicles (e.g. (30, 54)).

Finally, the scenario we modeled in this work bears a resemblance to other interactive scenarios. In essence, the scenario in Fig. 2A entails a continuous and dynamic interaction where participants search for a mutually beneficial solution. Although there are minor individual advantages to be gained regarding speed and comfort, the most important goal is mutual: to prevent a collision. These aspects are comparable to other interactive scenarios such as traffic interactions between vehicles and pedestrians (e.g. (45, 48)), pedestrian interactions in a crowd (55), or physical human-robot interactive tasks (56) (e.g. (50, 57)). The shortcomings of other models that led to the development of the CEI framework (10)—the assumption of rationality, absence of communication, and difficulties in extending game theory beyond high-level decisions—also apply to these related scenarios. Thus, exploring CEI-based models, such as the one presented here, for other interactive scenarios can be an interesting topic for future work.

Our work has three important limitations: the simplified scenario, the manually chosen model parameters, and the complexity of fitting an intermittent behavior model. To start with the first limitation, we have used a simplified merging scenario in this work to gain insight into the complex dynamics of driver interactions (38). This scenario enabled us to uncover the characteristics of human behavior regarding accelerations but does not include two important aspects of merging: steering and traffic rules (such as the right of way). We previously found that the characteristics of human behavior in our simplified simulator correspond to those found in real-world driving (13, 58, 59). Therefore, we are confident that our model captures an important aspect of human behavior in real traffic: intermittent piece-wise constant control. Since the other underlying principles of the model (i.e. risk perception (47), communication (11, 39), and a belief about intent (45)) have also been observed in real traffic and successfully used in other models, we believe that our model can be generalized to realistic merging scenarios. Nonetheless, extending the model to

realistic scenarios would require modifications such as using beliefs that cover multiple tactical responses to the same situation (60); it, therefore, remains a topic for future work.

Using the model for scenarios other than merging would also be possible. Scenarios where drivers can be reasonably assumed to use restricted (i.e. 1D) motions and only implicit communication, such as a pedestrian crossing in front of a vehicle, could be modeled with the current version of the model. However, the current model will not work (if unaltered) for scenarios with possibilities for multiple high-level actions (e.g. for lane-changing or intersections with left and right turns) or for scenarios where explicit communication (e.g. indicator lights) plays a major role.

Second, the model uses 10 manually chosen parameters to match the scenario (these were heuristically determined, based on literature, or tuned to fit the data; see the Methods section for details). Among these parameters are the length of the planning horizon for the drivers and the saturation time τ that governs how long it takes drivers to replan when the conflict is resolved. Although we found the model robust to changes in these parameters, how they generalize to other scenarios (e.g. with different dimensions of the track) is unknown. It is possible that all scenarios need a specific set of parameters and that to generalize the model to work in multiple scenarios simultaneously, these parameters need to be dynamically adjusted.

Finally, due to the intermittent nature of our model, it is complex to fit it to human data. Because of the mechanism where a risk threshold triggers a planning update, individual trials only provide limited information about the threshold value that would describe a driver best. If the driver acts, the threshold is exceeded, and if they don't act, the threshold is not exceeded. The amount of action is not related to the upper risk threshold. Therefore, we used a grid search method to obtain individual values for risk thresholds. However, this method is computationally inefficient, imprecise, and hard to use with more complex scenarios with multiple control inputs (i.e. when including steering). Because intermittent control is a key aspect of the CEI framework and our model, this is a fundamental limitation to the potential to generalize the model. More work is needed to develop a more robust method to fit intermittent models such as ours.

In conclusion, our model hypothesized that a communication-enabled kinematic belief combined with risk-based intermittent actions underlie human interactive behavior in merging. In contrast to the currently prevalent game-theoretic models of traffic interactions, our model does not rely on the assumption of rationality and explicitly includes implicit communication between drivers. Despite its simplicity, our model could accurately describe the joint behavior of human drivers and their individual contributions in merging interactions on three levels: control inputs, safety margins, and decisions. We believe our model could be a useful tool to increase the fundamental understanding of the effects of vehicle kinematics on the beliefs of other drivers. Therefore, we hope our model represents a step towards understanding driving interactions and developing interaction-aware automated driving.

Methods

In this work, we evaluate a CEI model in a simulated environment and compare it to human behavior data that was previously collected in a simulator experiment. Here we reiterate the details of the experiment, present the design of the four modules of the CEI model (plan, communication, belief, and risk perception), and discuss the parameters we used for the model and the fitting

procedure. Finally, we present the details of the software and data we used in this work which is all available online from public repositories.

Experiment and simulation environment

The data on human driver behavior we used in this work were previously gathered in an experiment in a coupled top-down-view driving simulator (13, 38). Eighteen volunteers (6 female, 12 male, mean age: 25, std: 2.6) participated in the study and were divided into nine fixed pairs (i.e. each participant interacted with the same counterpart in all trials). This experiment was approved by TU Delft's Human Research Ethics Committee and all participants gave their informed consent before participating.

The participants controlled the acceleration of their vehicle using the gas and brake pedal of a steering-wheel game controller (Logitech Driving Force GT). The headings of the vehicles were fixed (i.e. equal to the heading of the road). Participants each sat behind a computer screen that showed a top-down view of the simulation. They were seated in the same room behind a curtain to prevent them from seeing each other. The drivers wore noise-canceling headsets (Sony WH-1000XM3) with ambient music to prevent them from communicating in any other way than through vehicle kinematics.

In the experiment, the drivers followed a track consisting of three sections of equal length (50 m each, total track length 150 m): a tunnel, an approach and a car-following section. The vehicle dimensions were 4.5 m × 1.8 m. In the tunnel, the drivers could observe both vehicles and their initial velocity. The initial velocities were either equal (10 m/s) for both vehicles or one of the vehicles had a 0.8 m/s advantage (9.6–10.4 m/s). If the vehicles maintained their velocity, they would collide at the merge point with varying headways (i.e. distance from front bumper to front bumper). We called this the projected headway and varied it between 0, 2, and 4 m for each vehicle. Conditions were labeled according to the projected headway and relative velocity (e.g. -4.8) where positive numbers denote an advantage for the left vehicle. Drivers were told: "Maintain your initial velocity yet prevent a collision. No vehicle has the right of way. Remain seated, use one foot on the gas or brake pedal, keep both hands on the steering wheel, and do not communicate by making sounds or noise. Remember that this is a scientific experiment, not a game or a race."

The participants approached the merge point from the left or the right side (randomized before each trial). However, only their own view was varied; in the experimenter's view, and in all results discussed here, the same driver in a pair is referred to as the left or right driver. To facilitate velocity perception, the steering wheels provided vibration feedback when vehicles deviated from their initial velocity. This feedback increased with the magnitude of the deviation. In case of a collision, the simulation was paused for 20 s. This time penalty lasted longer than the duration of a single trial (~16 s), thereby providing an incentive to avoid collisions.

Both in the experiment and the model simulations, we simulated the vehicles as point mass objects, their dimensions were only used for collision detection. The vehicles were subject to a negative acceleration due to resistance of $a_r = 0.5 + 0.005v^2$, where v is the vehicle's velocity. Vehicle velocities were always positive.

Model

Our proposed model is based on the CEI modeling framework (10). According to this framework, joint driver behavior can be understood as a combination of four modules: plan, communication, belief, and risk perception. All of these modules are needed for

the model to function; none can be optionally left out (see the [Supplementary material](#) for an extended discussion).

Plan

The drivers are assumed to have a deterministic plan for the near future. The model plans to maintain a constant acceleration input over its planning horizon; the acceleration value is obtained by minimizing a cost function over this horizon (Eq. 1). This cost function J penalizes deviating from the desired velocity v_d (which in the experiment is equal to the initial velocity) and large values of the acceleration input a_t

$$J = \sum_{t=t_0}^{t_f} (v_t - v_d)^2 + a_t^2, \quad (1)$$

where v_t denotes the velocity at time t , and T is the time horizon. Importantly, the cost function does not include a term for collision avoidance; instead, the CEI framework assumes that drivers manage safety by keeping the risk below the risk threshold, which is imposed as an optimization constraint on the planning module of the model (10).

With the optimal constant acceleration, a trajectory (i.e. a set of waypoints over time) is constructed over the time horizon T . This trajectory is later used to evaluate the perceived risk. The planned constant acceleration is applied with added noise to execute the plan. This noise represents the discrepancy between a planned acceleration and the gas pedal input (i.e. inaccuracies in the driver's neuromuscular system and internal model). The noise is added after the optimization in the planning phase and remains constant until the next replan. Noise is drawn from a scaled normal distribution: $\mathcal{N}(\mu_n = 0, \sigma_n^2 = \frac{1}{40^2})$.

If the optimization fails because no solution can be found within the constraints, the model falls back to either full braking or full acceleration. If the ego vehicle is behind the other vehicle, heading for a collision, and no solution to the planning problem can be found, the model applies full braking. Similarly, if the ego vehicle is ahead of the other vehicle but cannot find a feasible plan, it applies full acceleration. In both cases, a new optimization is triggered at the next time step until a valid solution can be found again.

Communication

In our model, drivers communicate through vehicle kinematics. Explicit communication (e.g. with indicator lights) is not included in the model or the experiment for simplicity. Drivers observe the other vehicle's position, velocity, and acceleration at every time step. Position and acceleration observations are assumed to be perfect. Velocity perception is assumed to be noisy to account for the fact that drivers sometimes accelerate and sometimes decelerate at the tunnel exit in the same condition (Fig. 3A). This behavior can be explained by the fact that drivers under—or overestimate the other vehicle's velocity in the tunnel.

The noise in the velocity perception is inspired by evidence accumulation, a concept used in driver decision-making studies before (45, 61). Specifically, we assume that drivers update their perceived velocity of the other vehicle v^p at every time step with an observation affected by noise

$$v_t^p = v_{t-1}^p + dv^p \quad (2)$$

$$dv^p = \alpha(v_t^i - v_{t-1}^p) + \beta dW. \quad (3)$$

In Eqs. 2 and 3, subscript t denotes time, the superscript p denotes the perception of the other vehicle's velocity, dv^p is the perception

update, v^i is the other vehicles true (or real) velocity, α denotes the update rate, β the noise level, and W is a stochastic Wiener process (thus dW is a sample from a normal distribution $\mathcal{N}(\mu = 0, \sigma^2 = dt)$). The observed position, velocity, and acceleration (i.e. the incoming communication) are used to construct the belief.

Belief

The observed communication is used to create a belief about the future positions of the other vehicle. This belief consists of belief points at a specific belief frequency f_b over the same time horizon T as the one used in the plan. Every belief point is a probability distribution (represented by a probability density function) over positions. Each of these belief points are represented by the sum of two normal distributions:

$$b_t = \frac{1}{2} \mathcal{N}(\mu_t, \sigma_t^2) + \frac{1}{2} \mathcal{N}(\mu_t, \phi \sigma_t^2), \quad (4)$$

where b is the belief point representing the probability distribution over positions for the other vehicle at time t and ϕ is a scaling factor. The first part of this equation represents the positions of the other vehicle that are kinematically feasible within the bounds of comfortable acceleration. The second part is kinematically infeasible within comfortable bounds and can be interpreted as a belief that something unexpected will happen (e.g. an emergency braking). The motivation behind including two components to the belief distribution was that a single normal distribution either assigns similar (almost uniform) probabilities to the kinematically likely and unlikely positions (when σ_t is high) or only considers high-risk scenarios (when σ_t is low). Our belief model addresses this issue by emphasizing the kinematically likely outcomes, but at the same time including a safety margin in case of errors or unlikely events (e.g. emergency braking, long reaction times, or a perception error of the other driver).

The parameters μ_t and σ_t are based on a normally distributed expected acceleration ($a_t \sim \mathcal{N}(\mu_{a,t}, \sigma_{a,t}^2)$) which is constructed based on driver's memory about recent acceleration observations. Drivers keep a memory of recent acceleration observations of the other vehicle $M_{a,t}$

$$M_{a,t} = \{a_{-T_m}, \dots, a_t\} \quad (5)$$

$$\mu_{a,t} = \bar{M}_{a,t} \quad (6)$$

$$\sigma_{a,t}^2 = \left(\frac{1}{3} a_c\right)^2 + \text{var}(M_{a,t}). \quad (7)$$

Here, the mean expected acceleration $\mu_{a,t}$ is calculated as the average of the remembered values over the past T_m seconds (Eq. 5). The standard deviation $\sigma_{a,t}$ of the expected acceleration is based on the maximum comfortable acceleration a_c (assuming that 99.7% of observed accelerations fall within $\pm a_c$) and an added variance $\text{var}(M_{a,t})$. The latter part increases the expected variance in future accelerations if inconsistent behavior has recently been observed.

The parameters for the belief point distributions are then constructed using point mass kinematics with this normally distributed acceleration

$$\mu_t = \frac{1}{2} (t - t_0)^2 \mu_{a,t} + v_{t_0} (t - t_0) + p_{t_0}, \quad (8)$$

$$\sigma_t^2 = \frac{1}{2} (t - t_0)^2 \sigma_{a,t}^2, \quad (9)$$

Table 2. Manually designed model and simulation parameters.

T	6.0 s	dt	0.05 s
T_m	4.0 s	f_b	4 Hz
σ_n	$\frac{1}{40}$	β	0.6
τ	1.6 s	ϕ	3.0
α	0.5	a_c	$1.0 \frac{m}{s^2}$

where p_{t_0} denotes the observed position of the other vehicle, v_{t_0} is the perceived (noisy) velocity of the other vehicle, and t denotes the time of this belief point and t_0 the current time.

Risk perception

Risk perception combines the planned trajectory and the belief about the future positions of the other vehicle to calculate the probability of a collision. For every belief point, the model determines the bounds of collisions (10); these are the extremum positions of the other vehicle that will result in a collision. The believed probability that the other vehicles will be between these bounds is the perceived probability of a collision. This probability is assumed to be the perceived risk.

The perceived risk is evaluated against two dynamic risk thresholds, the upper (ρ_u) and the lower (ρ_l) threshold. Both thresholds consist of a driver's individual base value (θ), which is adjusted with an incentive function:

$$\rho_{u,t}^d = \theta_u^d + \lambda_{u,1}\Delta p_t + \lambda_{u,2}\Delta v_t + \lambda_{u,3}\Delta p_t\Delta v_t \quad (10)$$

$$\rho_{l,t}^d = \theta_l^d + \lambda_{l,1}\Delta p_t + \lambda_{l,2}\Delta v_t + \lambda_{l,3}\Delta p_t\Delta v_t. \quad (11)$$

In these equations, superscript d denotes a specific driver, Δp_t and Δv_t are the relative position and velocity from this driver's perspective at time t , and λ are the incentive parameters which are assumed to be constant over the population. We analyzed the differences between the model with these incentive function and a version with constant (static) risk thresholds in the [Supplementary material](#).

A replan is triggered if the upper risk threshold ρ_u is exceeded. This replan aims to find an acceleration that brings the perceived risk below $0.8\rho_l$. If the perceived risk stays below the lower threshold ρ_l for longer than the saturation time τ , the conflict is assumed to be resolved, and another re-plan is triggered to revert to "normal" behavior. In this case, the risk is constrained to be below $0.6\rho_u$. Finally, if the desired velocity is reached while the vehicle is accelerating or decelerating, another replan is triggered to allow the vehicle to maintain the preferred velocity.

Model assumptions

In summary, the model makes six main assumptions about human behavior; these were mentioned previously but are explicitly listed here for completeness. In line with the CEI-framework (10), which is based on the idea of satisficing (62), we assume drivers do not have the time or mental capacity to optimize their plan over multiple possible futures constantly. Therefore, we assume they use a single deterministic plan which they only update when the risk threshold is exceeded. Because of the track's simplicity and the top-down view, we assume drivers have full information about the road. Their perception is assumed to be noisy; however, reasonable amounts of noise in the position and acceleration observations have very little effect on the simulated behavior; thus, we assume only velocity perception is noisy. The experiment assumed that there is only implicit communication (38) for simplicity; we

Table 3. The fitted population-level parameters for the incentive functions.

Parameter	Value	Std. Err.	z	P
$\lambda_{u,1}$	0.003	0.001	3.27	0.001
$\lambda_{u,2}$	0.018	0.006	2.97	0.003
$\lambda_{u,3}$	-0.006	0.001	-6.38	0.000
$\lambda_{l,1}$	0.004	0.001	3.55	0.000
$\lambda_{l,2}$	0.016	0.008	1.95	0.051
$\lambda_{l,3}$	-0.003	0.001	-2.03	0.042

adopted the same assumption in the model. Furthermore, for the same reason of simplicity, we assume risk perception is equal between drivers while only their thresholds differ (for an extended discussion, see the discussion in Ref. (63)). Finally, the beliefs are assumed to be uni-model and centered around a constant velocity and acceleration prediction. This is motivated by the simplicity of the track as well. As discussed in the Discussion section, this final assumption should be re-evaluated when extending the model to more realistic scenarios.

Model parameter fitting

Our model and simulation use 10 parameters with values that were manually designed, their values are shown in Table 2. The timing parameters for the simulation (dt), planning (T), and belief (T_m, f_b) were chosen such that they are suitable for the scenario yet enable reasonable computation times. The noise parameters σ_n and β , and saturation time τ were manually tuned to reflect the human data. The belief scaling factor ϕ was designed to obtain sufficient resolution in the risk signals. The parameters α (which denotes how long drivers need to observe a vehicle to estimate its velocity) and a_c (the maximal comfortable acceleration) were based on empirical literature: α (64, 65), a_c (66).

The risk thresholds and incentive function parameters were fitted to the data using a grid search. We created a 25×25 grid for every kinematic condition using upper thresholds in the range [0.3, 0.9] and lower thresholds in the range [0.01, 0.4]. For these grids, we disabled the incentive functions and only used the base thresholds (θ in Eqs. 10 and 11). We ran one simulated trial per set of thresholds per condition with all the noise in the model disabled, resulting in 11 grids of 625 trials.

In every trial, we simulated the behavior of a single CEI driver against a vehicle traveling at constant velocity to obtain the immediate response of the CEI driver at the tunnel exit before any interaction takes place. For every such trial, we recorded the modeled driver's deviation from the initial velocity after 1.0 s.

Then, for every trial of every human participant, we searched the grid for the risk threshold values that best described that driver's velocity deviation after 1.0 s. This resulted in individual base threshold values (ρ_u and ρ_l) for every trial, in total 110 sets of thresholds for every participant. With 18 participants, this gave us 1980 upper and lower threshold values for 11 different kinematic conditions.

To determine participant-level thresholds based on these trial-level thresholds, we used two linear mixed-effect models, one for the lower and one for the upper threshold. Both models used the form: $\rho \sim \Delta p_t * \Delta v_t$, where ρ denotes the threshold and Δp_t and Δv_t are the relative position and relative velocity respectively. Random intercepts were included per participant. The resulting fixed-effect coefficients (Table 3) were used as incentive parameters (λ in Eqs. 10 and 11), and the random-effect coefficients (Table 1) were used as participant-level base values θ_l and θ_u .

Acknowledgments

We thank Lorenzo for his help with the Casadi implementation and Tom and Julian for their valuable feedback on the manuscript.

Supplementary Material

Supplementary material is available at PNAS Nexus online.

Funding

This research was partly funded by Nissan Motor Company, Ltd. and by the RVO grant TKI2012P01.

Author Contributions

Conceptualization: All authors. Methodology: All authors. Investigation: All authors. Formal analysis: All authors. Funding acquisition: D.A. Writing—original draft: O.S. Writing—review & editing: All authors.

Preprints

This manuscript was posted on a preprint server: <https://doi.org/10.48550/arXiv.2312.09776>.

Data Availability

We implemented a custom simulation environment in Python to run the experiment and simulate the model. The optimization in the planning part of the model was implemented with Casadi (67). The statistical analyses were performed in python using the statsmodels package (68). Our code is publicly available on Github (42). The data from the experiment (40) and model simulations (41) are available on the 4TU data repository. Interactive plots of the human data are available online (43).

References

- Clements LM, Kockelman KM. 2017. Economic effects of automated vehicles. *Trans Res Rec: J Trans Res Board*. 2606(1):106–114.
- Harper CD, Hendrickson CT, Mangones S, Samaras C. 2016. Estimating potential increases in travel with autonomous vehicles for the non-driving, elderly and people with travel-restrictive medical conditions. *Trans Res Part C: Emerging Technol*. 72:1–9.
- Pettigrew S. 2017. Why public health should embrace the autonomous car. *Aust N Z J Public Health*. 41(1):5–7.
- Brown B, Broth M, Vinkhuyzen E. 2023. The halting problem: video analysis of self-driving cars in traffic. In: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems.
- Krajewski R, Bock J, Kloeker L, Eckstein L. 2018 Nov. The highD dataset: a drone dataset of naturalistic vehicle trajectories on German highways for validation of highly automated driving systems. In: 2018 21st International Conference on Intelligent Transportation Systems (ITSC). Vol. 2018-November. IEEE. p. 2118–2125.
- Michon JA. 1985. A critical view of driver behavior models: what do we know, what should we do? In: *Human behavior and traffic safety*. US, Boston, MA: Springer. p. 485–524.
- Sadigh D, Landolfi N, Sastry SS, Seshia SA, Dragan AD. 2018. Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state. *Auton Robots*. 42(7):1405–1426.
- Schwartz W, Pierson A, Alonso-Mora J, Karaman S, Rus D. 2019. Social behavior for autonomous vehicles. *Proc Natl Acad Sci U S A*. 116(50):24972–24978.
- Siebinga O, Zgonnikov A, Abbink D. 2022. A human factors approach to validating driver models for interaction-aware automated vehicles. *ACM Trans Hum-Robot Interact*. 11(4):1–21.
- Siebinga O, Zgonnikov A, Abbink DA. 2023. Modelling communication-enabled traffic interactions. *R Soc Open Sci*. 10(5):230537.
- Brown B, Laurier E, Vinkhuyzen E. 2023. Designing motion: lessons for self-driving and robotic motion from human traffic interaction. *Proc ACM Hum-Comput Interact*. 7(GROUP):1–21.
- Kita H. 1999. A merging-giveway interaction model of cars in a merging section: a game theoretic analysis. *Trans Res Part A: Policy Pract*. 33(3–4):305–312.
- Siebinga O, Zgonnikov A, Abbink DA. 2024. Human merging behavior in a coupled driving simulator: how do we resolve conflicts? *IEEE Open J Int Trans Syst*. 5:103–114.
- Treiber M, Hennecke A, Helbing D. 2000. Congested traffic states in empirical observations and microscopic simulations. *Phys Rev E - Stat Phys Plasmas Fluids Relat Interdiscip Topics*. 62(2):1805–1824.
- Kauffmann N, Winkler F, Naujoks F, Vollrath M. 2018. “What makes a cooperative driver?” Identifying parameters of implicit and explicit forms of communication in a lane change scenario. *Trans Res Part F: Traffic Psychol Behav*. 58:1031–1042.
- Zhang Q, et al. 2020. A game theoretic model predictive controller with aggressiveness estimation for mandatory lane change. *IEEE Trans Int Veh*. 5(1):75–89.
- Choudhury CF. 2007. Modeling driving decision with latent plans [PhD dissertation] [PhD thesis]. Massachusetts Institute of Technology (MIT).
- Ahmed KI. 1999. Modeling drivers’ acceleration and lane changing behavior [PhD thesis]. Massachusetts Institute of Technology (MIT).
- Kondyli A, Elefteriadou L. 2011. Modeling driver behavior at freeway-ramp merges. *Transp Res Rec*. 2249:29–37.
- Laval JA, Leclercq L. 2008. Microscopic modeling of the relaxation phenomenon using a macroscopic lane-changing model. *Trans Res Part B: Methodol*. 42(6):511–522.
- Michaels RM, Fazio J. 1989. Driver behavior model of merging. *Transp Res Rec*. 1213:4–10.
- Hidas P. 2002. Modelling lane changing and merging in microscopic traffic simulation. *Trans Res Part C: Emerging Technol*. 10(5–6):351–371.
- Kesting A, Treiber M, Helbing D. 2007. General lane-changing model MOBIL for car-following models. *Transp Res Rec*. 1999: 86–94.
- Yang Q, Koutsopoulos HN. 1996. A microscopic traffic simulator for evaluation of dynamic traffic management systems. *Trans Res Part C: Emerging Technol*. 4(3 PART C):113–129.
- Wan X, Jin PJ, Yang F, Zhang J, Ran B. 2014. Modeling vehicle interactions during merge in congested weaving section of freeway ramp. *Trans Res Record J Trans Res Board*. 2421(1):82–92.
- Daamen W, Loot M, Hoogendoorn SP. 2010. Empirical analysis of merging behavior at freeway on-ramp. *Transp Res Rec*. 2188: 108–118.
- Marczak F, Daamen W, Buisson C. 2013. Merging behaviour: empirical comparison between two sites and new theory development. *Trans Res Part C: Emerging Technol*. 36:530–546.

- 28 Dong C, Dolan JM, Litkouhi B. 2018. Smooth behavioral estimation for ramp merging control in autonomous driving. *IEEE Int Veh Symp Proc.* 2018-June(Iv):1692–1697.
- 29 Coskun S, Zhang Q, Langari R. 2019 Jul. Receding horizon Markov game autonomous driving strategy. In: 2019 American Control Conference (ACC). Vol. 2019-July. IEEE. p. 1367–1374.
- 30 Li N, et al. 2018. Game theoretic modeling of driver and vehicle interactions for verification and validation of autonomous vehicle control systems. *IEEE Trans Control Syst Technol.* 26(5):1782–1797.
- 31 Liu H, Xin W, Adam Z, Ban J. 2007. A game theoretical approach for modelling merging and yielding behaviour at freeway on-ramp section. *Trans Traffic Theory.* 17(January):196–211.
- 32 Ji A, Levinson D. 2020. A review of game theory models of lane changing. *Trans A: Trans Sci.* 9935(May):1–19.
- 33 Dragan AD, Lee KCT, Srinivasa SS. 2013. Legibility and predictability of robot motion. *ACM/IEEE Int Conf Hum-Robot Interact.* 1: 301–308.
- 34 Guest O, Martin AE. 2021. How computational modeling can force theory building in psychological science. *Perspect Psychol Sci.* 16(4):789–802.
- 35 Brito B, Zhu H, Pan W, Alonso-Mora J. 2021. Social-VRNN: One-Shot Multi-modal Trajectory Prediction for Interacting Pedestrians. *Proceedings of the 2020 Conference on Robot Learning.* 155:862–872.
- 36 Mészáros A, Alonso-Mora J, Kober J. 2023. TrajFlow: learning the distribution over trajectories.
- 37 Salzmann T, Ivanovic B, Chakravarty P, Pavone M. 2020. Trajectron++: dynamically-feasible trajectory forecasting with heterogeneous data. *Computer Vision – ECCV 2020. ECCV 2020. Lecture Notes in Computer Science.* 12363:683–700. https://doi.org/10.1007/978-3-030-58523-5_40.
- 38 Siebinga O, Zgonnikov A, Abbink D. 2022. Interactive merging behavior in a coupled driving simulator: experimental framework and case study. *Hum Factors Trans.* 60:516–525.
- 39 Lee YM, et al. 2021. Road users rarely use explicit communication when interacting in today's traffic: implications for automated vehicles. *Cogn Technol Work.* 23(2):367–380.
- 40 Siebinga O, Zgonnikov A, Abbink DAD. 2022. Data underlying the publication: interactive merging behavior in a coupled driving simulator: experimental framework and case study [dataset].
- 41 Siebinga O, Zgonnikov A, Abbink DAD. 2023. Data underlying the publication: a model of dyadic merging interactions explains human drivers' behaviour from input signals to decisions [dataset].
- 42 Siebinga O. 2022. Simple-merging-experiment [Software] [accessed 2024 Sep 25]. <https://github.com/tud-hri/simple-merging-experiment>.
- 43 Siebinga O. 2024. Simple merging experiment, interactive online supplementary materials [accessed 2024 Sep 25]. <https://tud-hri.github.io/simple-merging-experiment/>.
- 44 Premack D, Woodruff G. 1978. Premack and Woodruff: chimpanzee theory of mind. *Behav Brain Sci.* 4:515–526.
- 45 Markkula G, et al. 2023. Explaining human interactions on the road by large-scale integration of computational psychological theory. *PNAS Nexus.* 2(6):1–13.
- 46 Tian R, Tomizuka M, Sun L. 2021 Sep. Learning human rewards by inferring their latent intelligence levels in multi-agent games: a theory-of-mind approach with application to driving data. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS); Prague, Czech Republic. IEEE. 4560–4567.
- 47 Kolekar S, de Winter J, Abbink D. 2020. Human-like driving behaviour emerges from a risk-based driver model. *Nat Commun.* 11(1):4850.
- 48 Camara F, Bellotto Nicola, Cosar Serhan, Weber Florian, Nathanael Dimitris, et al. 2021. Pedestrian models for autonomous driving part II: high-level models of human behaviour. *IEEE Trans Int Trans Syst.* 22(9):5453–5472. <https://doi.org/10.1109/TITS.2020.3006767>.
- 49 Gabler V, Stahl T, Huber G, Oguz O, Wollherr D. 2017. A game-theoretic approach for adaptive action selection in close proximity human-robot-collaboration. In: 2017 IEEE International Conference on Robotics and Automation (ICRA). p. 2897–2903. <https://doi.org/10.1109/ICRA.2017.7989336>.
- 50 Nikolaidis S, Nath S, Procaccia AD, Srinivasa S. 2017. Game-theoretic modeling of human adaptation in human-robot collaboration. *ACM/IEEE Int Conf Hum-Robot Interact. Part F1271:* 323–331.
- 51 Turnwald A, Olszowy W, Wollherr D, Buss M. 2014. Interactive navigation of humans from a game theoretic perspective. *IEEE Int Conf Int Robots Syst. (Iros):*703–708. <https://doi.org/10.1109/IROS.2014.6942635>.
- 52 Kolekar S, de Winter J, Abbink D. 2020. Which parts of the road guide obstacle avoidance? Quantifying the driver's risk field. *Appl Ergon.* 89(July):103196.
- 53 Ward E, Evestedt N, Axehill D, Folkesson J. 2017. Probabilistic model for interaction aware planning in merge scenarios. *IEEE Trans Int Veh.* 2(2):1–1.
- 54 Queiroz R, et al. 2022. A driver-vehicle model for ADS scenario-based testing. *IEEE Trans Intell Transp Syst.* 25(8):8641–8654. <https://doi.org/10.1109/TITS.2024.3373531>.
- 55 Martinez-Gil F, Lozano M, García-Fernández I, Fernández F. 2017. Modeling, evaluation, and scale on artificial pedestrians: a literature review. *ACM Comput Surv.* 50(5):1–35. <https://doi.org/10.1145/3117808>.
- 56 De Santis A, Siciliano B, De Luca A, Bicchi A. 2008. An atlas of physical human-robot interaction. *Mech Mach Theory.* 43(3): 253–270.
- 57 Sadrifaridpour B, Wang Y. 2018. Collaborative assembly in hybrid manufacturing cells: an integrated framework for human-robot interaction. *IEEE Trans Autom Sci Eng.* 15(3):1178–1192.
- 58 Markkula G. 2014. Modeling driver control behavior in both routine and near-accident driving. *Proc Hum Factors Ergonomics Soc.* 2014-Janua:879–883.
- 59 Markkula G, Boer E, Romano R, Merat N. 2018. Sustained sensorimotor control as intermittent decisions about prediction errors: computational framework and application to ground vehicle steering. *Biol Cybern.* 112(3):181–207.
- 60 Siebinga O, Zgonnikov A, Abbink D. 2023. Uncovering variability in human driving behavior through automatic extraction of similar traffic scenes from large naturalistic datasets. In: The 2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC 2023). IEEE.
- 61 Zgonnikov A, Abbink D, Markkula G. 2022. Should I stay or should I go? Cognitive modeling of left-turn gap acceptance decisions in human drivers. *Hum Factors.* 66(5):1399–1413. <https://doi.org/10.1177/00187208221144561>.
- 62 Simon HA. 1956. Rational choice and the structure of the environment. *Psychol Rev.* 63(2):129–138.

- 63 Siebinga O. 2024. Communication-enabled interactions in highway traffic: a joint driver model for merging [Dissertation] (tudelft), Delft University of Technology.
- 64 Fath AJ, Lind M, Bingham GP. 2018. Perception of time to contact of slow- and fast-moving objects using monocular and binocular motion information. *Atten Percept Psychophys*. 80(6): 1584–1590.
- 65 Jörges B, Harris LR. 2022. Object speed perception during lateral visual self-motion. *Atten Percept Psychophys*. 84(1):25–46.
- 66 Hoberock LL. 1977. A survey of longitudinal acceleration comfort studies in ground transportation vehicles. *J Dyn Syst Meas Control Trans ASME*. 99(2):76–84.
- 67 Andersson JAE, Gillis J, Horn G, Rawlings JB, Diehl M. 2018. CasADi – a software framework for nonlinear optimization and optimal control. *Math Programm Comput*. 11:1–36.
- 68 Seabold S, Perktold J. 2010. statsmodels: econometric and statistical modeling with python. In: 9th Python in Science Conference(SCIPY 2010).