



Delft University of Technology

Modelling communication-enabled traffic interactions

Siebinga, O.; Zgonnikov, A.; Abbink, D.A.

DOI

[10.1098/rsos.230537](https://doi.org/10.1098/rsos.230537)

Publication date

2023

Document Version

Final published version

Published in

Royal Society Open Science

Citation (APA)

Siebinga, O., Zgonnikov, A., & Abbink, D. A. (2023). Modelling communication-enabled traffic interactions. *Royal Society Open Science*, 10(5), Article 230537. <https://doi.org/10.1098/rsos.230537>

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.

Research



Cite this article: Siebinga O, Zgonnikov A, Abbink DA. 2023 Modelling communication-enabled traffic interactions. *R. Soc. Open Sci.* **10**: 230537.
<https://doi.org/10.1098/rsos.230537>

Received: 24 April 2023

Accepted: 26 April 2023

Subject Category:

Computer science and artificial intelligence

Subject Areas:

computer modelling and simulation

Keywords:

driving interactions, driver modelling, traffic, communication

Author for correspondence:

O. Siebinga

e-mail: o.siebinga@tudelft.nl

Electronic supplementary material is available online at <https://doi.org/10.6084/m9.figshare.c.6632547>.

Modelling communication-enabled traffic interactions

O. Siebinga, A. Zgonnikov and D. A. Abbink

Department of Cognitive Robotics, Delft University of Technology Mekelweg 2, Delft, The Netherlands

OS, 0000-0002-5614-1262

A major challenge for autonomous vehicles is handling interactions with human-driven vehicles—for example, in highway merging. A better understanding and computational modelling of human interactive behaviour could help address this challenge. However, existing modelling approaches predominantly neglect communication between drivers and assume that one modelled driver in the interaction responds to the other, but does not actively influence their behaviour. Here, we argue that addressing these two limitations is crucial for the accurate modelling of interactions. We propose a new computational framework addressing these limitations. Similar to game-theoretic approaches, we model a joint interactive system rather than an isolated driver who only responds to their environment. Contrary to game theory, our framework explicitly incorporates communication between two drivers and bounded rationality in each driver's behaviours. We demonstrate our model's potential in a simplified merging scenario of two vehicles, illustrating that it generates plausible interactive behaviour (e.g. aggressive and conservative merging). Furthermore, human-like gap-keeping behaviour emerged in a car-following scenario directly from risk perception without the explicit implementation of time or distance gaps in the model's decision-making. These results suggest that our framework is a promising approach to interaction modelling that can support the development of interaction-aware autonomous vehicles.

1. Introduction

Autonomous and automated vehicles (AVs) hold the potential to help address major societal challenges related to mobility and sustainability. However, one of the major open problems in autonomous vehicle development is safely and acceptably dealing with driving scenarios that require *reciprocal* interaction with human road users. In these interactions, such as in highway merging or intersection negotiation, both vehicles influence and respond to the actions of each other. It entails quick and sometimes iterative negotiations, based on communication (see

e.g. [1–3]) that can either be implicit (vehicle motions) or explicit (e.g. honking, signalling). The continuous dynamics of a reciprocal interaction govern safety, priority (who goes first, who gives way), and acceptance (by passengers and other road users). For example, drivers can be misunderstood or cause annoyance by being too conservative or aggressive (interfering with or ignoring others' communication). Therefore, fundamental knowledge about continuous human reciprocal interactions is necessary to develop and evaluate safe and acceptable AV behaviour for these scenarios. However, this fundamental knowledge about the dynamics of interactions is currently lacking. We advocate using a modelling approach for human reciprocal traffic interactions to develop the fundamental understanding that in the future can help design better AV behaviour.

Modelling is a common way of gaining an understanding of human driving behaviour. But it has so far mostly been done with a focus on single-driver behaviour, either in single-vehicle (e.g. [4,5]) or multi-vehicle scenarios such as car following [6,7], lane changing [8,9], and gap acceptance [10,11]. Most multi-vehicle approaches assume that the modelled driver responds to other traffic participants, but that they don't respond in turn. For example, car-following models assume that the following driver responds to the leading vehicle, but this leading vehicle does not change its behaviour based on the follower's actions. We call this the *one-way interaction* assumption because the influence on behaviour is unidirectional. This assumption disentangles the behaviours of the multiple drivers and thereby enables the researchers to better understand and model the behaviour of the driver of interest. The scope of these models is thus deliberately restricted to a single driver. This one-way interaction assumption is justified for car-following models and the likes, but not for interactive driving scenarios such as merging or intersection negotiations, which are inherently reciprocal. Simply joining two one-way interaction models to describe an interaction will neglect the drivers' beliefs about the other's future actions and their expected influence on them. Furthermore, it also neglects the presence and effects of communication between the drivers. Therefore, we argue that the scope of an interaction model should include all participants to begin with. It should be a model of a joint interactive system.

The current mainstream approach to modelling joint interactive systems in traffic (as opposed to individual drivers) is by using game theory. Game theory was developed as a framework to describe reciprocal interactions between players in abstract games. It has been used extensively to model traffic interactions. The first model of human merging behaviour based on game theory was proposed in 1999 by Kita [12]. In 2007, Liu *et al.* improved the game theoretical approach by removing the assumption of constant velocity [13]. After that, many works followed (e.g. [14–17]). However, applying game theory to model dynamics between two drivers is not trivial, because game theory makes three strong assumptions about these players.

First, there is the assumption that all players rationally maximize some utility function. Empirical evidence has shown that even in simple economic games [18], but also in driving behaviour [19] and traffic interactions [20], this assumption does not hold for human players. Second, game theory does not allow communication between the players—an aspect known to be important in interactive driving scenarios [3]. Third, the majority of game-theory-based interaction models use a set of discrete actions for the drivers. Although this is useful to describe the higher-level tactical [21] decisions of drivers accurately (for example, the decision to yield or merge), it does not describe the lower-level operational [21] dynamics of the interaction (e.g. changes in velocity or trajectory). Therefore, these approaches are not sufficiently detailed for developing safe and acceptable AV behaviour. Combined, these three limitations motivate the need for an alternative approach to modelling reciprocal traffic interactions that allows for communication, bounded rationality, and continuous dynamic actions.

To address this gap, we propose a framework for communication-enabled-interaction (CEI) modelling. It can be used to create model implementations, of which we provide one example in a case study.¹ Our modelling framework relaxes the common (game-theoretic) assumptions that drivers are rational agents and have full information about the strategies of other drivers. It is based on the notion that all drivers have a plan they want to execute and a belief about what other drivers are going to do. Combined, this plan and belief result in a perceived risk for every driver. The drivers are assumed to act to keep this risk below their individual threshold. The key insight of the framework is that the beliefs about others are updated based on communication between the agents. In a simulation case study, we show that the implementation of a CEI model produces plausible behaviour of two interacting drivers in a simplified merging scenario. Besides that, human-like gap-keeping behaviour emerges directly from the notion of

¹The software implementation of the presented model and its simulation environment are available online at [22]. The data discussed in the results section can be found at [23].

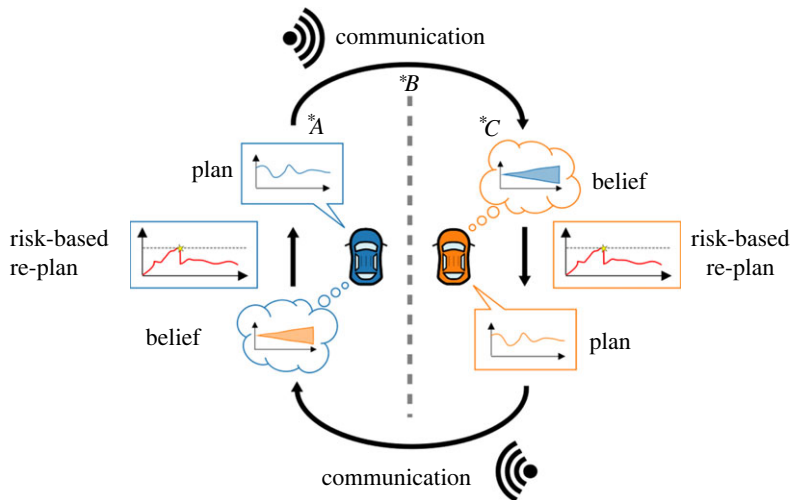


Figure 1. An overview of the proposed communication-enabled-interaction (CEI) modelling framework. This framework is designed to capture the reciprocal interaction between two drivers, rather than the one-way interaction behaviour of one driver with respect to another. Each driver has a *plan* for their own behaviour. Plan updates are triggered based on a *risk threshold* and a *risk estimate* arising from a *belief* of how the other driver will move over time. Each driver *communicates* their plan (intention) either implicitly (e.g. through vehicle motion) or explicitly (e.g. through light signals) to the other driver. This communication links one driver's plans to the belief of the other and can be divided into three components denoted $*A$, $*B$ and $*C$. $*A$ represents the mapping of a driver's plan to its communication, $*B$ represents the means of communication and $*C$ denotes the belief update of the other driver based on the received communication.

risk perception. These results show that the proposed modelling framework provides a promising new approach for modelling human–human driver interactions.

2. Communication-enabled-interaction modelling

We propose a framework to model reciprocal human–human traffic interactions between two drivers. This framework captures the joint system of both drivers rather than a single driver responding to their environment. The framework explicitly includes (implicit/explicit) communication between the drivers, which facilitates the joint interaction. Each of the two drivers is described by four components: a notion of that driver's perceived risk, a deterministic plan for the driver's own control behaviour (e.g. accelerating/decelerating), a means of communication, and a probabilistic belief about the future behaviour of the other driver (figure 1). The general framework we present here only defines loose requirements for how these four components should be operationalized. When implementing the model for a specific scenario or use case, these components can be operationalized based on existing literature (e.g. from the fields of human behaviour modelling, traffic communication, intent inference, or vehicle path planning). This means the model framework allows the incorporation of different methods to operationalize each of the four components, without having to fully redesign it. In this section, we will discuss the four components and our reasoning behind their functionality and requirements. The assumptions and requirements that need to be taken into account when implementing a model based on this framework will also be discussed per component. In §3, we will illustrate how each component can be implemented in an example implementation for a simplified merging scenario.

2.1. Framework components

2.1.1. Risk-based re-plan

Recent research in non-interactive scenarios has shown that a quantification of perceived risk can explain driver adaptations in speed and lateral position and can be used to accurately predict future trajectories [4,24,25]. In our framework, we combine this notion of risk-based decision-making in driving with Simon's ideas of *bounded rationality* [26] and *satisficing* [27]. *Bounded rationality* implies that humans are not capable of fully optimizing their behaviour all the time. *Satisficing* (a portmanteau of *satisfy* and *suffice*) is an example of bounded rationality in which humans are assumed to not continuously search for an optimal solution. Instead, they are *satisfied* with a 'good enough' solution that *suffices*. We

reason that the only solutions that suffice and satisfy in a driving interaction are the ones that are *subjectively safe enough*. To formalize these ideas and combine them in a framework, we hypothesize that drivers act to keep their perceived risk below their personal risk-threshold.

Using such a threshold incorporates Simon's ideas in two ways. First, it defines what solutions are subjectively safe enough. Second, it limits (or bounds) the cognitive capacities (or effort) required from the driver because it allows the driver to only rethink their plan when the situation has changed and the current plan does not suffice or satisfy them anymore. This is what we call a *risk-based re-plan* (figure 1). By incorporating these ideas, we step away from the fundamental assumption of game theory that humans are rational utility maximizers and move towards a formulation that allows for team effort and mutual goals.

In summary, our framework assumes every driver evaluates the risk of their current deterministic plan, given their probabilistic belief about what other drivers are planning to do. Risk perception can be based on a number of factors, such as high velocity, high acceleration, or the probability of a collision. This evaluation happens continuously, but drivers will only perform a re-plan if the perceived risk exceeds their threshold. This should result in drivers with a low risk-threshold adapting their plan in an early stage of the interaction to reduce the estimated risk. At the same time, drivers with a high risk-threshold will instead continue their current plan and take advantage of the fact that the risk of the situation is lowered by the other driver. Intuitively, this can be explained as the driver with the higher risk-threshold being more aggressive.

2.1.2. Plan

The second component in our framework is the *plan*. We assume that drivers have a deterministic plan about the actions they will take in the immediate future. In the framework, this plan takes the form of a deterministic set of waypoints over a limited time horizon. This time horizon should be long enough to include (part of) the interaction.

The construction of this plan (i.e. the planning algorithm) should only consider features that are not related to risk and safety (e.g. desired velocity or comfort), as the perceived risk is constantly evaluated separately to determine if the current plan still suffices and satisfies. This evaluation is done taking into account both the plan and the belief. When re-planning, the risk threshold should be used as a constraint in the planning algorithm. As long as such a constraint can be imposed, the plan can be constructed using any suitable path-planning algorithm.

2.1.3. Communication

One of the key concepts of the framework is that drivers actively communicate their plan to other drivers. This assumption is based on field studies on human–human traffic interactions that confirm that traffic participants actively communicate their plan both explicitly and implicitly to others (e.g. [3]). Experiments on other (non-driving) tasks that require team effort have shown that humans use their movement actions to coordinate with their team member [28] (which is a form of implicit communication). The assumption of communication can also be effectively used to model human behaviour in those tasks [29]. Finally, in simulation, communication can be beneficial for controlling co-bots that navigate among humans [30], resulting in fewer dead-lock situations. In summary, previous research suggests that humans communicate in traffic and that the assumption of communication can be used for both the effective modelling of human teamwork behaviour and the effective control of robots.

In the CEI modelling framework, communication links the *plan* of one driver to the *belief* of the other driver. In practice, this means that three aspects of communication need to be designed when implementing a CEI model. First, one needs to determine the mode of communication; What signals are used to communicate? These signals can be explicit (e.g. turn indicators) or implicit (e.g. velocity, heading angle, or acceleration). Second, a mapping from a plan to its communication is required. This can be as simple as just executing the plan, but one could come up with more elaborate mappings based on traffic communication studies—for example, slowing down purely to communicate that the other driver can go first (for an example of modelling such exaggerated trajectories in a bottle-grasping task, see [29]). Finally, a mapping from communication to belief is needed; this mapping specifies how a probabilistic belief is updated based on the received communication.

2.1.4. Belief

Both drivers are assumed to have probabilistic beliefs about what the other driver will do in the near future. This belief consists of a number of points over a time horizon. Each of these belief points is represented by a




Figure 2. A top-down view of the simplified merging scenario as used in the case study, rotated 90° clockwise. Vehicles follow predefined paths (road centres) that merge at a pre-defined merge point. Vehicles have a two-dimensional body (4.5 m × 1.8 m) and their headings change instantly at the merge point. The model controls the accelerations of the vehicles directly. The dimensions of the track are defined by two parameters. Distance l_a (25 m) denotes the distance between the start points of the vehicles. Distance l_b (50 m) is the distance to travel from the start point until the merge point, and from the merge point until the end of the track.

probability distribution over positions for the other driver for that specific time in the future (figure 1). This assumption is based on the intuition that human drivers have a general but uncertain idea about what other drivers are planning to do, which is a concept that has been successfully applied in other modelling frameworks such as belief-desire-intention programming (based on [31]) and (Bayesian) theory of mind [32].

When implementing the belief part of the CEI model, the only requirement is that the chosen probability distribution can be updated using new information (coming from the observed communication). In practice, this means that most parametric probability distributions are suitable because they can be updated with methods such as Bayesian updates.

3. Case study: an example of an implementation

To demonstrate the feasibility of the proposed model framework and to investigate the effects of design choices (parameters) on model behaviour, we have implemented a CEI model for a simplified merging scenario. In this case study, we show that even with simple components the model framework can produce plausible, human-like interactive behaviour. At the same time, it is not the purpose of this case study to quantitatively assess the model's consistency with human behaviour. Such an assessment using fine-grained data on the interactive behaviour of two drivers requires a detailed investigation and is therefore left for future work.

3.1. Simplified merging scenario

For this case study, we used a simplified symmetric merging scenario (figure 2). In this scenario, two vehicles approach a merge point on a predefined track. The model can directly control the acceleration of the vehicles, but there is no steering involved. The vehicles have a rectangular bounding box for collision detection. The heading of the vehicles is predefined and always corresponds to the heading of the road. At the merge point, the heading of the vehicles changes instantly.

The vehicles in the simplified scenario are subject to a negative acceleration due to resistance and drag. The net acceleration (a^{net}) is the applied input (a^{in}) minus the negative acceleration a^r (a function of the vehicle's velocity v):

$$a^{net}(v) = a^{in} - a^r(v), \quad (3.1)$$

where

$$a^r(v) = \alpha v^2 + \beta. \quad (3.2)$$

Parameters α and β define the magnitude of the drag and constant resistance, respectively ($\alpha = 0.0005$ and $\beta = 0.1$). Besides the resistance, the vehicles have a maximum acceleration $a^{max} = 2.5 \text{ m s}^{-2}$, which is




Figure 3. An example to illustrate the plan and the belief of the model. Panel (a) shows four (of 80) deterministic plan points along the one-dimensional track. These are the planned centre positions of the own vehicle at four points in time. The distributions represent the believed centre position of the other vehicle at the same four (of 16) points in time, where colours denote the points in time. Panel (b) shows these plan and belief points after a single belief update. This update increased the certainty of the belief about the other vehicle's position. The belief is updated at every time step. Panel (c) shows the risk evaluation for one of the points. To evaluate the risk, the probability of a collision (p_c) is evaluated by calculating the probability that the other vehicle will be within the bounds of collision for the given planned position. This risk evaluation is done at every time step for all belief points. If the maximum perceived risk value exceeds the upper risk-threshold, a re-plan is triggered. This re-plan uses the perceived risk as a constraint for the optimization. To lower the risk, the planned position could be moved in the direction of the black arrow.

the same for positive and negative accelerations. The velocity of the vehicles is restricted to non-negative values. The simulation updates all dynamics at a rate of 20 Hz.

3.2. Plan

The planning part of the model consists of a path-planning algorithm that minimizes the following cost function:

$$c = \sum_{n=1}^N (v_n - v^d)^2 + (a_n^{\text{in}})^2, \quad (3.3)$$

where n denotes the time step and v the vehicle's velocity. This cost function includes terms for minimizing the squared input a^{in} and for travelling at a desired velocity v^d . The path is planned at the same frequency as the simulation (20 Hz) and is subject to a time horizon of 4 s ($N=4/0.05=80$).

A visual example of the plan, belief, and risk perception is shown in figure 3. When initially planning the path, the cost function of equation (3.3) is minimized, so an optimal path is found with respect to comfort and speed (figure 3a). If, at the next time step, the current plan still satisfies (i.e. the risk threshold is not exceeded), the current plan is continued. We assume that maintaining velocity at the final time step is the practical equivalent of maintaining the current plan.

When the risk threshold is exceeded, the cost function is minimized again to find a new plan (figure 3c). This time, the minimization is subject to a risk constraint. Based on the idea of satisficing, we hypothesize that humans do not spend unlimited effort to find an optimal plan, but instead search for a new solution that satisfies and suffices. We hypothesize that re-planning is easiest (i.e. requires the least cognitive effort) if the new plan is close to the previous plan (i.e. uses the same strategy). Therefore, the re-planning optimization is executed with the old plan as the initial condition. When using a gradient-descent algorithm, this will result in a solution that is close to the previous plan while the risk constraint is met. For example, if the current plan is to decelerate and pass behind the

other driver, the most likely outcome of the re-planning will be to decelerate even more and increase the gap. This will lower the perceived risk while using the current strategy.

If the optimization with the current plan as the initial condition does not succeed, three other initial conditions are considered: full braking at all time steps, no acceleration input at all time steps, and full acceleration at all time steps. The candidate plan with the lowest cost is used as the initial condition for a second re-plan. This can result in a change of strategy, but only if the current strategy is not feasible anymore. For example, when the driver was decelerating but decelerating even more will not reduce the risk enough, the driver will investigate if acceleration will reduce the risk and change its strategy if needed.

3.3. Belief

The belief is kept as a sequence of probability distributions over positions for the other vehicle, each at a specific point in time (figure 3a). This sequence of belief points uses the same time horizon as the planning part of the model (4 s) but contains fewer points for simplicity. Belief points are kept at a 4 Hz frequency (this number was based on an initial evaluation of the model), resulting in a sequence of $4 \times 4 = 16$ points. Each belief point is represented by a Gaussian distribution.

The Gaussian distributions are initialized by combining the initial velocity and position of the other vehicle with the maximum bounds of acceleration. To initialize a belief point, the mean of the Gaussian is set to the position that corresponds to the other driver maintaining their current velocity. To calculate the standard deviation, upper and lower position bounds (ub and lb) are used. These are calculated by predicting the position of the other vehicle if it applies the maximum and minimum possible acceleration continuously. The standard deviation is then calculated as the difference between the bounds and the mean divided by 3 ($\sigma = (ub - \mu)/3$). The factor one-third is based on the fact that 99.73% of the area under a normal distribution corresponds to $\mu \pm 3\sigma$. Once the simulation time is equal to the timestamp corresponding to the first belief point, this point is removed from the sequence and a new point is initialized.

3.4. Communication

Human communication during driving is a complex topic on which a lot of research has been done. Thus, there is much potential for including complex communication models based on empirical evidence in a CEI model. However, for this initial investigation of the modelling framework, we used a simple implicit communication model that does not include any explicit communication signals (e.g. turn indicators). We only use velocity and position as communication signals. These two values are assumed to be constantly observed by the other driver without any errors or noise.

When sending communication, the drivers do not use a mapping from their current plan to the actions they take. Instead, they just take the next action from their plan. When receiving communication, drivers use a constant velocity model combined with bounds of comfortable acceleration to update their belief. All belief points are updated at every time step using Bayesian updating.

3.4.1. Updating the belief

For Bayesian updating, the previous belief point serves as the prior distribution, and the resulting posterior is adopted as the updated belief point (figure 3b). The likelihood is constructed using the constant-velocity model. We assume the likelihood to be a Gaussian distribution where the standard deviation is constant and known. This means the likelihood and prior form a conjugate pair, meaning that the posterior will also be a Gaussian distribution of which the μ and σ^2 have a closed-form solution. The likelihood function for the belief point at time t is defined as follows:

$$\mathcal{N}\left(\mu = \frac{p}{t}, \sigma^2 = \left(\frac{a_c t}{6}\right)^2\right). \quad (3.4)$$

In this equation, p denotes a position sampled from the prior (the previous belief point), t denotes the time corresponding to the belief point, and a_c is the maximum comfortable acceleration ($a_c = 1.0 \text{ m s}^{-1}$). The same value is used for positive and negative accelerations, thus the distribution is symmetrical. The likelihood function describes the probability of observing a velocity v (now), given a sampled predicted position p (at time t) from the prior belief. The mean μ corresponds to constant velocity, and σ is determined based on the assumption that 99.73% of the distribution falls within the bounds of comfortable acceleration.

With this likelihood function, the posterior has a closed-form solution. We denote the prior as $\mathcal{N}(\mu_0, \sigma_0^2)$ and the posterior as $\mathcal{N}(\mu_1, \sigma_1^2)$. When updating with a single data point v , the solution for the posterior becomes²

$$\mu_1 = \frac{\mu_0 \sigma^2 + v \sigma^2 (1/t)}{\sigma^2 + \sigma_0^2 (1/t^2)} \quad (3.5)$$

and

$$\sigma_1^2 = \frac{\sigma^2 \sigma_0^2}{\sigma^2 + \sigma_0^2 (1/t^2)}. \quad (3.6)$$

3.5. Risk

The risk perceived by the drivers is assumed to be proportional to the probability of a collision. Other aspects (i.e. high velocity and high acceleration) are assumed not to contribute to the perceived risk for simplicity. To estimate the probability of a collision, we define the concept of *bounds of collision* (figure 3c). These are the extreme positions of the other vehicle that would result in a collision, given the position of the own vehicle. These bounds are calculated for every point in the driver's plan. For example, if we know the driver will be at position x at time t , we can use the vehicles' dimensions to calculate that a collision will occur if and only if the other vehicle is at a position between $x + c_1$ and $x - c_2$ at the same time; these are the bounds of collision. The believed probability that the other vehicle will be within these bounds at that time can be calculated using the belief about the other vehicle's position. This probability is then equal to the probability of a collision at that time.

The perceived risk for a complete plan is determined by taking the maximum risk over all belief points. A re-plan is triggered if the perceived risk exceeds an upper threshold ρ_u . Only using the upper threshold, however, poses a potential problem when the merging conflict is resolved because after that there will be no triggers to re-plan anymore. This might cause vehicles to stall or drive very slowly for no reason. We avoid this by extending the risk module with a lower risk-threshold ρ_l and a saturation time τ . If the perceived risk is lower than ρ_l and the last update was longer than τ ago, a re-plan is also triggered. When a re-plan optimization is performed, the perceived risk is constrained to be lower than the average of the two thresholds. For the implementation of this constraint, the instant heading change at the merge point in the track posed a problem. Therefore, a linear approximation of the bounds of collision is used.

3.6. Investigated scenarios

In total, every driver in the model has four parameters that determine their behaviour: a desired velocity v_d , an upper risk-threshold ρ_u , a lower risk-threshold ρ_l , and a saturation time τ . Besides these parameters, the initial velocity and position (v_0 and x_0 respectively) of the drivers can also be adjusted. Both drivers always start from the beginning of the track. In the case study, we investigate the effect of these parameters and the effect of differences in the initial condition in four scenarios (table 1).

The first two scenarios (A and B) manipulate the initial and desired velocities of the right driver while keeping the parameters of the left driver fixed; the drivers here have the same risk thresholds. In scenario A, the drivers are not expected to be on a collision course if they stick to their desired velocity, but in scenario B, they are.

Scenarios C and D focus on the risk thresholds. Scenario C investigates the effect of a difference in risk thresholds between drivers. Scenario D investigates the sensitivity of model behaviour to variations of these thresholds in one of the drivers. The saturation time τ only affects the behaviour after the conflict is resolved, therefore it is kept constant at 2.0 s for all scenarios.

4. Results

4.1. Scenario A: no expected collision

Scenario A serves as a baseline scenario. Here, both drivers have an initial velocity that is equal to their desired velocity, but that differs from the velocity of the other driver (table 1). If they keep their initial (desired) velocity up until the merge point, no collision would occur. The left driver would pass the

²For a complete derivation of this closed-form solution, see the electronic supplementary material [33].




Figure 4. Model behaviour in scenario A (no expected collision). Line colours correspond to the vehicle colours in figure 2. (a) Positions of the left and right vehicles over time. The x positions of the vehicles are plotted with an offset to prevent the lines from overlapping after the merge point. The grey dots and dashed lines indicate vehicle positions at equal time stamps with an interval of 1.0 s. (b) Velocities of the vehicles over time. The stars indicate the moment when the simulated driver performed a re-plan because the upper risk-threshold was exceeded, and a circle denotes a re-plan because the risk fell below the lower threshold. These re-plans are only triggered if the last re-plan was longer than τ ago. (c) Accelerations of the vehicles over time. (d) Perceived risk of both simulated drivers. In case of a re-plan, the perceived risk after the re-plan is shown. The dashed horizontal lines in the lowest plot indicate the risk thresholds of the drivers. In this scenario, the drivers increased the small projected gap, even though they were initially not on a collision course. The simulated drivers behaved in a way to increase the initially narrow safety margin.

Table 1. Parameters of the investigated scenarios. Underlined values denote deviations from the default values. ρ_l and ρ_u denote the lower and upper risk thresholds, respectively. v_0 and v_d are the initial and desired velocities, respectively. x_0 denotes the initial position of the vehicle along the track.

	side	ρ_l	ρ_u	v_0	v_d	x_0
units	—	—	—	m s ⁻¹	m s ⁻¹	m
Condition A: no expected collision	left	0.2	0.5	10.0	10.0	0.0
	right	0.2	0.5	<u>9.0</u>	<u>9.0</u>	0.0
Condition B: on a collision course	left	0.2	0.5	10.0	10.0	0.0
	right	0.2	0.5	<u>9.0</u>	<u>9.0</u>	<u>1.2</u>
Condition C: high and low thresholds	left	0.2	<u>0.4</u>	10.0	10.0	0.0
	right	<u>0.3</u>	<u>0.6</u>	10.0	10.0	0.0
Condition D: threshold sensitivity	left	<u>0.3</u>	<u>0.4</u>	10.0	10.0	0.0
	right	<u>0.3</u>	<u>0.6</u>	10.0	10.0	0.0




Figure 5. Model behaviour in scenario B. The simulated drivers prevent a collision by slowing down. Initially, they both slow down, but after approximately 1.0 s, the left (initially faster) driver speeds up and reaches the merge point first. For details of the notation, see the caption of figure 4.

merge point first with a small distance gap of 0.2 m. Therefore, we would expect a rational optimizing model (that does not explicitly include human-like gap-keeping) to maintain the desired velocity all the way. A behaviour expected from human drivers, on the other hand, is to increase this small safety margin. In an empirical study [34], it was found that human drivers in the Netherlands merged on three different highway locations with mean headways of 12.6, 13.4, and 36.1 m for velocities below $60 \text{ km h}^{-1} = 16.7 \text{ m s}^{-1}$, and standard deviations of, respectively, 10.3, 12.8, and 18.2 (the headway is defined as the gap plus the leading vehicle length).

In the modelled outcome of scenario A (figure 4), the left driver reached the merge point first. They accelerated slightly to increase the safety margin at the merge point; after that, they returned to their desired velocity. The headway when the second vehicle reached the merge point was 6.4 m. This corresponds to the expected human behaviour and can not be modelled with utility-maximization unless utility is explicitly awarded for keeping a gap. The right driver did not take any action in this scenario. The reason for that is highlighted in the risk perception plot. The left driver's risk increases earlier because they expect to reach the merge point earlier. This increase causes the left driver to take action to lower the risk, while the right driver can continue their plan without exceeding their risk threshold. The right driver's perceived risk also decreases as soon as the left driver takes action; the right driver perceives that the conflict was resolved by the left driver.

4.2. Scenario B: on a collision course

In scenario B, the drivers have the same desired and initial velocities as in scenario A. However, the right vehicle starts with a 1.2 m head start. Therefore, the projected positions of the two vehicles at the merge point overlap by 1.0 m. Thus, if neither driver deviates from their desired velocity, this scenario will result in a collision. We would therefore expect that this scenario requires more severe action to be resolved than scenario A, but we do expect the model to avoid a collision.




Figure 6. Model behaviour in scenario C. The right driver maintains their initial velocity longer. After briefly decelerating, they accelerate and reach the merge point first. For details of the notation, see the caption of figure 4.

The modelled outcome of scenario B (figure 5) shows that this scenario indeed requires more effort from both drivers to resolve the conflict compared with scenario A. Both drivers start braking until the left driver decides they can only reduce the risk of a collision by accelerating. This can be explained by the fact that the left driver has a slightly higher velocity at this point compared with the right driver. The right driver sticks to their plan and keeps decelerating until the risk drops below the lower threshold and the saturation time has passed; only then do they accelerate again. This behaviour results in a safety margin between the vehicles that is not explicitly included in the reward function. Because the left driver is the first to accelerate, they reach the merge point first. This explainable interactive behaviour combined with the collision-free outcome can be regarded as a plausible human-like interaction.

4.3. Summary of scenarios A and B

In scenario A, the driver with the higher desired velocity that approached the merge point first also passed the merge point first. But the distance gap between the vehicles was enlarged by the drivers. This corresponds to what we expect from human drivers. If the drivers approach the merge point with an expected collision (scenario B), however, the drivers take more drastic action but still manage to resolve the conflict by interacting with each other.

4.4. Scenario C: high and low thresholds

Scenario C represents a case where the simulated drivers of both vehicles have the same initial conditions and desired velocities, but different risk thresholds. Compared with the previous scenarios, the right driver has a higher risk threshold while the left driver has a lower threshold. The left driver, having a lower threshold, is expected to act early in the interaction to reduce their perceived risk. In terms of human behaviour, this corresponds to risk-averse, conservative driving. The right driver (high

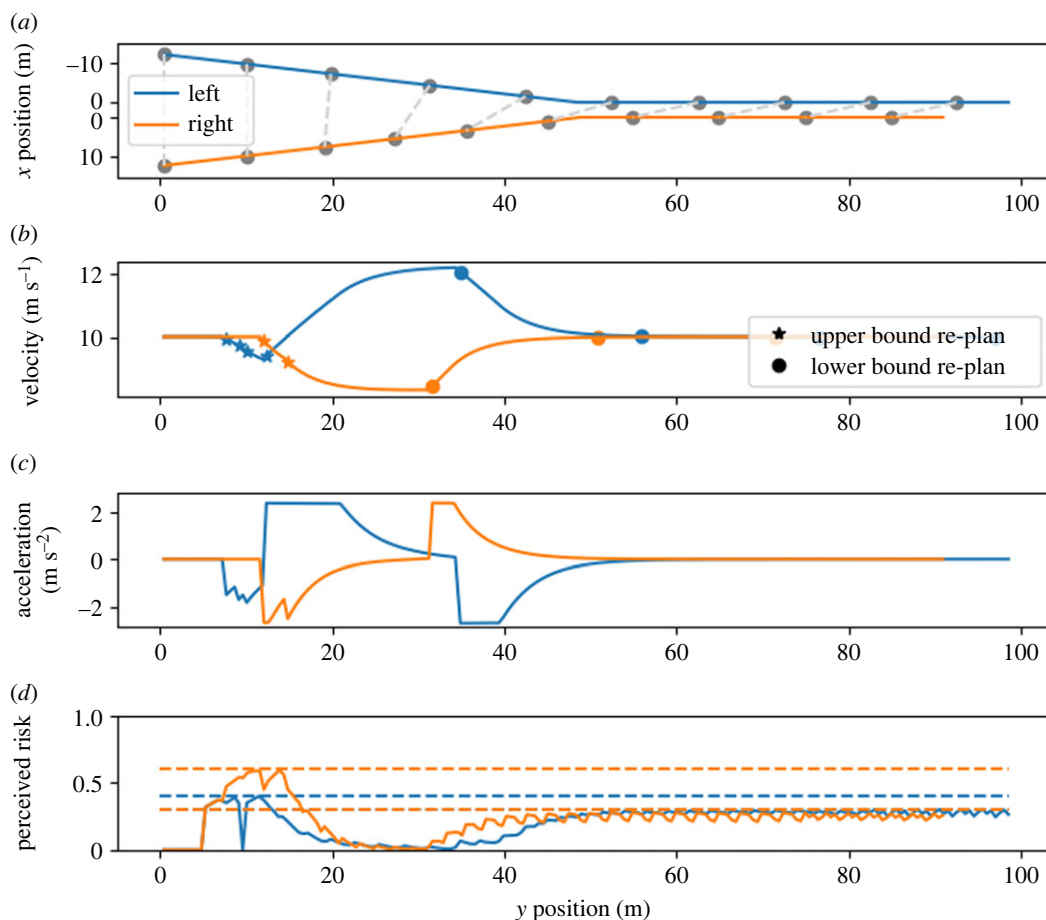


Figure 7. Model behaviour in scenario D. The slight change in ρ_l for the left driver (in comparison to scenario C) resulted in a major change in high-level outcome. Instead of the right driver, the left driver now reaches the merge point first. For details of the notation, see the caption of figure 4.

threshold, meaning higher tolerance to risk) is expected to react to a potential conflict at a later point and therefore to keep their velocity at the desired level longer. We expect that the right driver reaches the merge point first, and deviates less from their desired velocity compared with the left driver.

The modelled outcome of scenario C (figure 6) is as expected: the left driver reached their upper threshold first and started to decelerate to reduce the perceived risk. In terms of human driving, this can be seen as more conservative behaviour. The right driver reacts later because their risk threshold is exceeded at a later moment. They briefly decelerate, but quickly start to accelerate to reduce the risk since the left driver already decelerated. This results in the right driver reaching the merge point first and deviating less from their desired velocity than the left driver. This corresponds to the intuition that lower sensitivity to risk (i.e. higher risk thresholds) could be associated with more aggressive behaviour.

4.5. Scenario D: threshold sensitivity

Scenario D investigates the sensitivity of the modelled drivers' behaviour to variations in the lower risk-threshold. This scenario is similar to scenario C, except that the left driver has a slightly higher value for ρ_l (lower risk-threshold). We, therefore, expect very similar outcomes in scenarios C and D. The only expected difference is that the left driver in scenario D re-plans more frequently because the risk for the new plan is constrained to the average of the two risk thresholds. With a smaller difference between ρ_l and ρ_u , the absolute risk decrease at the re-plan points is smaller. This should cause the perceived risk to reach the upper threshold quicker and thus result in more frequent re-plan events.

However, the model simulation results show major differences between scenarios C and D (figures 6 and 7). As expected, the smaller difference between the left driver's lower and upper risk-thresholds resulted in more plan updates. But unexpectedly, this more frequent re-planning resulted in the left

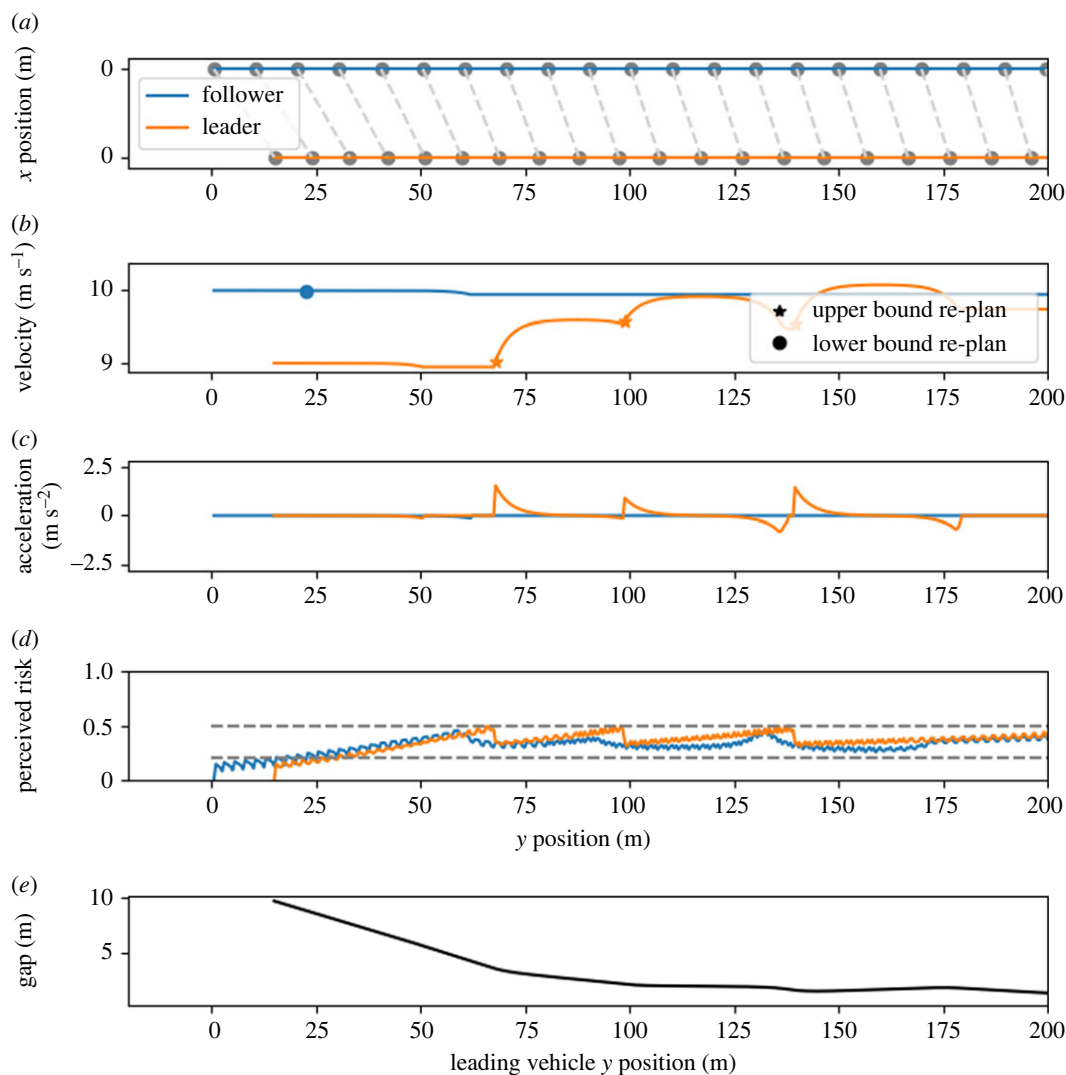


Figure 8. Model behaviour in the straight road scenario. For details of the notation, see the caption of figure 4. The bottom panel shows the gap between the vehicles as a function of the leading vehicle's position. In this scenario, the blue (following) vehicle has a higher desired velocity compared with the orange (leading) vehicle. The x -axes have been cropped to the first 200 m of the 400 m track.

driver starting to accelerate and reaching the merge point first. To keep their perceived risk under control, the left driver deviated from their desired velocity to a larger extent than the right driver. This observation can be explained by the fact that high velocities and accelerations do not contribute to risk. The left driver takes whatever action is needed to keep the probability of a collision below their threshold (in this case, high acceleration and high velocity). The slight change in risk thresholds and more frequent re-plans resulted in one of the re-plans initially failing. This triggered a change in the left driver's high-level strategy—they accelerated instead of braked—and this heavily influenced the outcome.

4.6. Summary of scenarios C and D

In scenario C, the driver with the higher risk-threshold (the right driver) passed the merge point first. This driver changed their plan at a later moment compared with the other driver. In terms of human behaviour, this can be explained as being more aggressive. The effect of slight changes to the lower threshold was shown to be substantial in scenario D. A small change resulted in a different interaction strategy, making the theoretically more 'conservative' left driver arrive at the intersection first. This more conservative driver used high velocities and accelerations to lower their perceived risk, even though high velocities would be interpreted by many human drivers as high-risk behaviour. The reason for this seemingly counterintuitive model behaviour is that the high velocities and accelerations on their own do not contribute to the perceived risk of these modelled drivers.




Figure 9. Steady-state gap sizes (averaged over the last second) on a straight road where the following vehicle has a higher desired velocity. The velocity difference between the vehicle is 10% and the initial time gap is 1 s.

4.7. Emergent gap-keeping behaviour for car following

Although the main focus of our model is on the interactive behaviour of the drivers when approaching the merge point, it also provides insight into their behaviour after the merging conflict is resolved. Specifically, in the four scenarios above, we found that the simulated drivers continued maintaining a gap on the straight section after the merge point. This behaviour was not explicitly programmed and the planner has no cost associated with short time or small distance gaps (a feature frequently used in human driver models [35,36]). Instead, these distance gaps appear to emerge from the combination of risk perception and a probabilistic belief about the plan of the other driver.

To further examine this effect, we investigated a scenario without a merge point. In this scenario, the drivers drive behind each other on a straight stretch of road (400 m). We used the default parameters from table 1, except for the velocity parameters. The leading vehicle has lower desired and initial velocities (9 m s^{-1}) compared with the following vehicle (10 m s^{-1}). Figure 8 shows that a steady-state gap emerges after approximately 100 m. In this scenario, the leading driver mostly acts to reduce the risk and prevent a collision.

Although the fact that the leading, not the following, driver mostly acts to maintain this gap is not uncommon for human drivers and has been observed under some conditions [37], it is not the most common behaviour for reducing the risk during car following [38]. We identified two causes for this model behaviour. First, the belief and risk perception in the model are purely symmetrical. There is no difference in perceived risk between drivers that are in front of or behind another driver, nor is there any difference in believed probability that the other driver will accelerate or decelerate. In natural traffic, this simplification will not hold. This should be accounted for when extending the model for use in those scenarios. Second, the risk thresholds of both drivers are equal in this example. It can be expected that in other situations, even under the previously mentioned assumption, the driver with the lower risk-threshold will act to maintain the gap, as was seen in scenario C. This can be either the leading or the following driver, as was observed in human behaviour [37,38].

We investigated the effect of absolute velocities on the resulting steady-state distance gap, where we took the average gap over the final second of simulation as the steady-state gap. We simulated the model behaviour in this scenario for different velocities, each time with a 10% velocity difference between the drivers, and an initial time gap of 1 s. We found that the emerging steady-state gap increased linearly with increasing velocities (figure 9). This corresponds to human behaviour: the same linear relationship has been previously observed in a study on human gap-keeping behaviour on highways at low speeds [39].

Our model explains this relationship between velocity and distance gap as follows. The leading driver (orange) is unsure about the future plan of the following driver (blue). It is possible that the blue driver will accelerate in the near future; in this case, a collision might occur. Because the orange driver keeps their risk below threshold, they will keep a distance from the blue driver to make sure that their own plan does not overlap too much with the possible future positions of the blue driver. Higher

velocities, with the same maximum comfortable acceleration, result in a high standard deviation in the belief points. This causes the gap size to increase with velocity.

The mentioned study [39] also showed that humans keep larger gaps (approx. 12 m to 23 m for the same velocity range) compared with our model. We, therefore, conclude that the model qualitatively captures the underlying risk-mitigation mechanism in human car-following behaviour, but needs to be further explored to investigate if fitting the model parameters to human data would also allow it to capture the magnitude of the gap characteristic of human drivers.

5. Discussion

In this work, we have proposed a modelling framework for reciprocal human–human interactions in traffic. We illustrated the utility of the framework by implementing a concrete model based on the framework, targeted at interactive behaviour in a simplified merging situation. We investigated the model's behaviour in four scenarios: one where the drivers are not on a collision course, one where they are, and two where we investigated the effects of the model parameters. The model captures the actions of two drivers who (i) successfully resolve merging conflicts without collisions, (ii) increase safety margins that are clearly too small (a 20 cm gap) for human drivers, and (iii) exhibit individual conservative and aggressive behaviour, based on physically meaningful model parameters: their risk thresholds. In all scenarios, the model behaves in a plausible way that corresponds to intuitions about human interactive behaviour in merging conflicts.

Furthermore, from the model's underlying principle (the notion of risk combined with the probabilistic belief about the other driver's plan), plausible behaviour emerged outside of the situations we developed and tuned the model for. Specifically, a realistic gap-keeping behaviour emerged, where the drivers kept larger distance gaps at higher velocities, as humans do [39]. This behaviour was observed even though no distance or time gap-related costs are incorporated in the model. These results show that the proposed model framework is a promising novel approach for modelling reciprocal multi-agent interactions in traffic.

Modelling interactions in traffic has both practical and fundamental applications. In practice, a modelling framework like the one we propose could aid the development of autonomous vehicle controllers that aim to increase acceptability and safety in interactive scenarios. More fundamentally, such modelling, even when limited to an isolated traffic scenario, could contribute to gaining fundamental knowledge of human behaviour by highlighting the cognitive mechanisms humans use when interacting with each other. Our novel framework addresses the limitations of existing modelling and control approaches, among which are game-theoretic models and interaction-aware controllers, because it explicitly incorporates communication and reciprocal interaction. Furthermore, our model framework does not make strong assumptions about human behaviour, such as the assumption that humans are rational utility maximizers. We hope that the initial exploration of the model framework presented here can spark a new strain of interaction modelling research.

5.1. Similar approaches

Among existing approaches to modelling traffic interactions, by far the most explored one is game theory. For example, for an extensive review of game-theory-based lane-changing models, see [40]. Similar to our framework, game theory aims at modelling joint interactive systems instead of modelling only one driver responding to another (for examples, see [12–17]). But, the difference is that our approach is not limited by two main assumptions (rationality and lack of communication) and—for the majority of game-theoretic approaches—a focus on decision-making without describing operational behaviour. Finally, and more conceptually, game-theoretic models implicitly approach traffic interactions as a competition, while in our framework the agents have a joint primary objective (interaction safety) that makes the interaction a cooperative effort.

In contrast with game theory, our approach explicitly incorporates the driver's ability to communicate their plan to other drivers, implicitly or explicitly. Although there are similarities with game theory—for example, our case study uses the same modality of communication as many game-theoretical approaches, namely position and velocity observations (e.g. [13,41]; for an overview, see [40])—there are two fundamental distinctions in how we approach communication.

First, the communication in our framework allows drivers to construct and update a belief about the other driver's plan without the need for any prior information about the other driver. This is a fundamental contradiction with game theory, where players are assumed to know each other's utility functions (at least

partially) beforehand. Therefore, in game theory, communication is not necessary because players can reason about what the other player is going to do to maximize their utility given the current state. Thus, the observations of position and velocity only serve to determine the state of the world. While in our model, position and velocity are used to convey information about the intentions of drivers.

Second, in game theory, observations are not ‘remembered’. They only serve to determine the *current* state, which is enough to reason about the other players’ actions. Previous states are irrelevant. This is also known as the Markov condition or assumption. While in our work, the history of communication is kept in the belief about the other driver’s intentions. Thus, the belief about a driver’s future actions is based on their recent behaviour, not only on the current state. Some approaches combine game theory with an online estimation of the other player’s utility function, thereby indirectly basing the belief about future actions (which directly depends on the utility function) on recent behaviour (e.g. [36,42]). However, in these approaches, the conveyed information is not regarded as intentional communication. Furthermore, these approaches only estimate part (e.g. a single parameter) of the utility function online; the rest is assumed to be known *a priori*.

Another modelling concept that bears resemblance to our approach is that of *belief-desire-intent (BDI) modelling*. BDI modelling is based on the philosophical work of Bratman [31] and models single agents that have a belief, a desire (goal), and an intent (plan). Many implementations of BDI models have been proposed for different applications [43]. The BDI framework and our CEI framework share the concepts that agents construct a (probabilistic) belief about other agents and the world, and then make a plan based on that belief to reach a final goal. The BDI framework, however, was not intended to account for interactions. It is primarily a model framework for individual agents that perform individual tasks. It therefore also does not incorporate communication but instead updates its beliefs based on changes that occurred in the world.

Finally, an important concept that can be complementary to the CEI-model framework, and bears resemblance to the BDI framework, is the concept of *Theory of Mind (ToM)* [44] (for examples of applications to human–robot interaction, see [45,46]). ToM is a psychological concept that assumes humans have an internal model of the beliefs, goals, and intentions of other humans in an interaction. Thereby, this means having the ability to reason about what other humans want and how they will try to achieve that goal. This idea that humans understand the mechanisms behind the actions and beliefs of others could be used in an implementation of our proposed CEI-model framework, which, in principle, only requires humans to form a basic belief about the future movements of others. As an example, the implementation of the CEI model in the case study assumes drivers predict where the other driver is going, not why they are doing that. A complete ToM model could extend this belief about future actions of the other, with a model of that driver’s beliefs and goals. Implementing a CEI-based model with an internal ToM model is an interesting avenue for future research.

Besides these different types of modelling approaches, recently a great deal of effort was put into approaches for controlling (autonomous) vehicles in merging scenarios (e.g. [36,42,47]). Although the underlying techniques (such as finding a policy by optimizing some utility function) are similar, the goal of these approaches is very different. While modelling approaches (such as ours) aim to best describe human behaviour, control approaches aim to find a safe and optimal solution to a control problem. Game theory can therefore be very suitable for use in control approaches (as was done in [36,42,48]).

Two recent works on modelling come close in scope to this work. In 2022, Markkula *et al.* proposed a modelling approach for individual agents in a driver–pedestrian interaction rather than multiple agents in a driver–driver interaction [49]. Using different versions of a model that incorporates a variety of concepts from psychology, with varying levels of complexity, they conclude that ‘modelling of human road user interaction is a formidable challenge’ [49]. Similar to our work, their findings suggest that the problem cannot be solved with simple rational models. Besides that, accounting for specific, previously unexplained, phenomena observed in human interactive behaviour could only be done using complex cognitive models. These conclusions resonate with our argument that the development of new model frameworks that go beyond game theory and the assumption of one-way interaction is a necessary step to improve our understanding of human traffic interactions.

Secondly, in 2014, Wan *et al.* also proposed an approach to model vehicle–vehicle interactions on merging ramps [50]. As in our work, they specifically address the reciprocal influence vehicles have on each other. Their (and our) work, therefore, differs from traditional driver models that usually describe a single driver responding to—but not influencing—other traffic. Another similarity between our proposed framework and the work by Wan *et al.* is that we both explicitly consider communication between vehicles. However, the model proposed by Wan *et al.* specifically targets

congested traffic and uses different mathematical models for vehicles that have different roles in the interaction (i.e. they determine who will lead, follow, and merge *a priori*). Wan *et al.* also do not consider individual differences between drivers.

5.2. Framework extensions

Although we have only demonstrated our proposed model framework for a simple merging scenario with two vehicles, it could easily be extended to more vehicles or to traffic interactions with other types of participants. The underlying reason is that while we put the model's bounding box around the complete interaction, the drivers within the model are strictly separated; the only component connecting the two drivers is communication (figure 1). This has two main advantages. First, communication in our framework is based on observable signals (e.g. turn indicators or velocity). This means that sending and receiving communication can easily be shared between multiple drivers, i.e. the communication is broadcast to all surrounding road users rather than sent directly to one of them. For that reason, the model framework can be extended to any number of drivers without requiring a redesign. Second, because the drivers are separated, it is possible to swap one of the drivers in the model with another type of agent—for example, a pedestrian. This would require adding the agent type to the observed communication, but since this is also an observable feature, it would not make the model more complex.

One could even go as far as replacing one of the agents in the model with a non-model agent altogether. This could, for example, be used to let a real human interact with the model in a driving simulator (this would require an optimized model implementation capable of running in real-time). This in turn would allow for the possibility of human drivers subjectively evaluating the ability of the model to describe natural interactions. Alternatively, a model could be used to evaluate autonomous vehicle controllers by letting the model interact with such a controller. Another potential extension useful for AV development is integrating the model into an AV controller to help it make decisions with an online evaluation of potential outcomes of an interaction.

We believe our model could also be adapted to other types of human–human interaction tasks. An example of such a task is cooperative bottle reaching, for which a communication model was developed in [29]. The task in [29] is similar to our task in that it constitutes a joint effort for which communication and action take place along the same channel (velocity/acceleration in our case). The main difference between our model framework and the communication model in [29] is that we target the interaction dynamics, in which we assume communication plays an important role, instead of targeting to model the communication as a stand-alone feature.

5.3. Limitations and future work

Both the specific model implementation and the general modelling framework have important limitations. To start with the former, the model used for the simplified merging scenario uses very simplistic implementations for all components. The plan is based on desired velocity and acceleration alone. The beliefs are one-dimensional and assumed to be Gaussian distributions. The communication is assumed to be perfect (continuous without any noise) and only based on implicit cues. And finally, the risk is only based on collision avoidance, not influenced by high velocities or accelerations. In future implementations of the model, these limitations need to be addressed and more realistic (and complex) model components should be investigated. However, it is important to first identify which of these limitations (if any) play a role in the model's ability to accurately reproduce human–human interactions. This could be done by comparing the model to data on human–human interactions gathered in a driving simulator experiment.

Another limitation of the current model implementation lies in the updates of the belief function. The assumption that the likelihood function (used for the Bayesian updates) has a known and fixed standard deviation results in the fact that every update reduces the standard deviation of the posterior, even if the new information contradicts the current belief. This is counterintuitive: contradictory information (incoming through communication) should increase the variability of the belief, not decrease it. Put differently, if another person or driver sends unclear communication about what they are going to do by alternating between accelerating and braking, one should keep all options open, not decrease the standard deviation of the predicted position after a couple of seconds while shifting the mean around on every time step. How to properly address this limitation remains an open question.

Finally, the model's satisficing-based decision-making can result in unstable outcomes for high-conflict scenarios. When re-planning, the drivers in the model will first search for a new solution close to the previous solution. For example, if the previous plan was to brake, the driver will first explore if braking harder will satisfy the new constraint. Only if this optimization fails, the driver will explore other strategies (i.e. acceleration) to lower the perceived risk. This drastic change in high-level behaviour is thus triggered by the first optimization failing. Therefore, slight numerical or temporal differences in this optimization can lead to different high-level outcomes, especially for situations that are highly symmetrical (e.g. when drivers have very similar parameters and none of the vehicles has a clear kinematic advantage). This was already observed in scenario D, where a slight change in model parameters caused a different outcome, but a similar outcome change could also result from changes in the type of numerical optimization solver or its parameters. One way of addressing this sensitivity is to make the model stochastic: introducing variability in the model's behaviour will make the outcome in high-conflict scenarios inherently stochastic and therefore could help to make it less sensitive to small external perturbations.

Adding stochasticity also addresses the main limitation of the overall framework, which is that currently, the framework is fully deterministic: with the exact same parameters (for model and solver), the model will always produce the same behaviour. This is inconsistent with the substantial behavioural variability that humans exhibit in traffic [51]. We see multiple possible ways of introducing stochasticity in the framework to account for this. To name two: adding stochasticity could be done in the receiving of communication (translating perceptual information to an updated belief) by using evidence accumulation mechanisms [10] or additive noise, or by including noise directly in the risk perception. However, more work is needed to determine the best approach.

A second limitation of the overall framework concerns improvements and redesigns of the model. Although the different components in the framework are separated, which should allow for easy redesign of parts of the model, they do depend on each other. This could mean that when redesigning one aspect of the model, a redesign of another aspect is inevitable. As an example, in the case study, we used velocity and position as the means of communication. These values are directly used in the belief update. However, if we change the communication component of the model, the belief and its update also need to be changed. This is an important consideration when starting a redesign of the model since this could be the case for more components.

Finally, event-based triggering of the re-plan based on perceived risk results in an uneven computational requirement from the model: some time steps may take significantly more time to compute than others. A result of this is that our current implementation of the model cannot run in real-time. Instead, we used offline simulation for the case study. This could pose a problem when an experiment needs to be performed where the model interacts directly with a human.

Although the presented case study shows promising results, there is much future work to be done on the proposed framework. In addition to accounting for stochasticity in human behaviour and optimizing the runtime performance of the model, a necessary next step is to compare the model to human–human interactive behaviour. However, even validating single-driver models that do not incorporate interactions is already a complex task [52], and therefore comparing our model to human–human interaction data requires a separate detailed investigation.

6. Conclusion

In this paper, we proposed a novel framework to model human–human driving interactions. The key insight underlying this framework is the focus on the joint behaviour of the drivers during the interaction, rather than the isolated behaviour of a single driver. The framework explicitly includes communication between drivers and mutual influences (reciprocal interaction). We implemented the model for a simplified merging scenario and investigated its behaviour in four scenarios. We conclude the following:

- The model avoids impending collisions via plausible driver–driver interactive behaviours;
- Changing the risk-threshold parameters per driver results in changes in behaviour that can be interpreted as more aggressive or conservative;
- Velocity-dependent gap-keeping behaviour emerges from the combination of risk-based planning and a probabilistic belief about other drivers' plans. With this behaviour, the model shows a fundamental aspect of human driving behaviour, without it being explicitly programmed;
- The proposed model framework is a promising novel approach for modelling two-way multi-agent interactions in traffic.

Data accessibility. The data underlying this publication have been uploaded to the 4TU data repository [23]. The software with the simulation environment and model implementation is available on GitHub [22] (as a maintained software repository) and on the 4TU data repository [53].

The mathematical derivation of the belief update function is provided in electronic supplementary material [33].

Authors' contributions. O.S.: conceptualization, data curation, formal analysis, investigation, methodology, project administration, software, validation, visualization, writing—original draft, writing—review and editing; A.Z.: conceptualization, methodology, project administration, resources, supervision, visualization, writing—review and editing; D.A.A.: conceptualization, funding acquisition, methodology, project administration, resources, supervision, visualization, writing—review and editing.

All authors gave final approval for publication and agreed to be held accountable for the work performed therein.

Conflict of interest declaration. We declare we have no competing interests.

Funding. This research was partly funded by Nissan Motor Company, Ltd and by the RVO grant no. TKI2012P01.

References

- Potzy J, Feuerbach M, Bengler K. 2019 Communication strategies for automated merging in dense traffic. *IEEE Intelligent Vehicles Symp., Proc.*, vol. 2019, pp. 2291–2298. Paris, France. 09–12 June 2019. IEEE. (doi:10.1109/IVS.2019.8813835).
- Ulbrich S, Grossjohann S, Appelt C, Homeier K, Rieken J, Maurer M. 2015 Structuring cooperative behavior planning implementations for automated driving. *IEEE Conf. on Intelligent Transportation Systems, Proc., ITSC*, vol. 2015, pp. 2159–2165. Gran Canaria, Spain. 15–18 September 2015. IEEE. (doi:10.1109/ITSC.2015.349).
- 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**, 367–380. (doi:10.1007/s10111-020-00635-y)
- Kolekar S, de Winter J, Abbink D. 2020 Human-like driving behaviour emerges from a risk-based driver model. *Nat. Commun.* **11**, 4850. (doi:10.1038/s41467-020-18353-4)
- Salvucci DD, Gray R. 2004 A two-point visual control model of steering. *Perception* **33**, 1233–1248. (doi:10.1068/p5343)
- Treiber M, Hennecke A, Helbing D. 2000 Congested traffic states in empirical observations and microscopic simulations. *Phys. Rev. E* **62**, 1805–1824. (doi:10.1103/PhysRevE.62.1805)
- Kesting A, Treiber M, Helbing D. 2010 Enhanced intelligent driver model to access the impact of driving strategies on traffic capacity. *Phil. Trans. R. Soc. A* **368**, 4585–4605. (doi:10.1098/rsta.2010.0084)
- Rahman M, Chowdhury M, Xie Y, He Y. 2013 Review of microscopic lane-changing models and future research opportunities. *IEEE Trans. Intell. Transp. Syst.* **14**, 1942–1956. (doi:10.1109/TITS.2013.2272074)
- Salvucci DD, Liu A. 2002 The time course of a lane change: driver control and eye-movement behavior. *Transp. Res. Part F: Traffic Psychol. Behav.* **5**, 123–132. (doi:10.1016/S1369-8478(02)00011-6)
- Zgonnikov A, Abbink D, Markkula G. 2020 Should I stay or should I go? Cognitive modeling of left-turn gap acceptance decisions in human drivers. *Hum. Factors* **0**(0). (doi:10.1177/00187208221144561)
- Tian K, Markkula G, Wei C, Lee YM, Madigan R, Merat N, Romano R. 2022 Explaining unsafe pedestrian road crossing behaviours using a psychophysics-based gap acceptance model. *Saf. Sci.* **154**, 105837. (doi:10.1016/j.ssci.2022.105837)
- Kita H. 1999 A merging-giveway interaction model of cars in a merging section: a game theoretic analysis. *Transp. Res. Part A: Policy Practice* **33**, 305–312. (doi:10.1016/S0965-8564(98)00039-1)
- Liu H, Xin W, Adam Z, Ban J. 2007 A game theoretical approach for modelling merging and yielding behaviour at freeway on-ramp section, pp. 196–211. Transportation and Traffic Theory. London, England 23–25 July 2007. Elsevier.
- Meng F, Su J, Liu C, Chen W-H. 2016 Dynamic decision making in lane change: Game theory with receding horizon. In *2016 UKACC 11th Int. Conf. on Control (CONTROL)*, pp. 1–6. Belfast, UK. 31 August–02 September 2016. IEEE. (doi:10.1109/CONTROL.2016.7737643)
- Tian R, Li S, Li N, Kolmanovsky I, Girard A, Yildiz Y. 2018 Adaptive game-theoretic decision making for autonomous vehicle control at roundabouts. In *2018 IEEE Conf. on Decision and Control (CDC)*, vol. 2018-December, pp. 321–326. IEEE Miami, FL, USA. 17–19 December 2018. (doi:10.1109/CDC.2018.8619275)
- Zhang Q, Filev D, Tseng HE, Szwabowski S, Langari R. 2018 Addressing mandatory lane change problem with game theoretic model predictive control and fuzzy Markov chain. In *2018 Annual American Control Conference (ACC)*, vol. 2018-June, pp. 4764–4771. IEEE Milwaukee, WI, USA. 27–29 June 2018. (doi:10.23919/ACC.2018.8431530)
- Coskun S, Zhang Q, Langari R. 2019 Receding horizon Markov game autonomous driving strategy. In *2019 American Control Conf. (ACC)*, vol. 2019-July, pp. 1367–1374. IEEE Philadelphia, PA, USA. 10–12 July 2019. (doi:10.23919/ACC.2019.8815251)
- Camerer CF. 2003 *Behavioral game theory: experiments in strategic interaction*. Princeton University Press.
- Schmidt-Daffy M. 2014 Prospect balancing theory: bounded rationality of drivers' speed choice. *Accid. Anal. Prevent.* **63**, 49–64. (doi:10.1016/j.aap.2013.10.028)
- Kalantari AH, Yang Y, Merat N, Markkula G. 2023 Modelling vehicle-pedestrian interactions at unsignalised locations: road users may not play the Nash equilibrium. *PsyArXiv* 1–34. (doi:10.31234/osf.io/axseu)
- Michon JA. 1985 a critical view of driver behavior models: what do we know, what should we do? In *Human Behavior and Traffic Safety*, pp. 485–524, Boston, MA, USA: Springer.
- Siebinga O. 2022 CEI-model repository. See <https://github.com/tud-hri/cei-model>.
- Siebinga O, Zgonnikov A, Abbink DA. 2022 Data underlying the publication: modelling communication-enabled traffic interactions (Version 1) [Data set]. 4TU.ResearchData. (<https://doi.org/10.4121/20749069.V1>)
- 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**, 103196. (doi:10.1016/j.apergo.2020.103196)
- Jensen E, Luster M, Yoon H, Pitts B, Sankaranarayanan S. 2022 Mathematical models of human drivers using artificial risk fields. In *IEEE Conf. on Intelligent Transportation Systems, Proceedings, ITSC*, vol. 2022-October, pp. 909–916. Macau, China 08–12 October 2022. (doi:10.1109/ITSC5140.2022.9922389)
- Simon HA. 1955 A behavioral model of rational choice. *Q. J. Econ.* **69**, 99. (doi:10.2307/1884852)
- Simon HA. 1956 Rational choice and the structure of the environment. *Psychol. Rev.* **63**, 129–138. (doi:10.1037/h0042769)
- Sacheli LM, Tidoni E, Pavone EF, Aglioti SM, Candidi M. 2013 Kinematics fingerprints of leader and follower role-taking during cooperative joint actions. *Exp. Brain Res.* **226**, 473–486. (doi:10.1007/s00221-013-3459-7)
- Pezullo G, Donnarumma F, Dindo H. 2013 Human sensorimotor communication: a theory of signaling in online social interactions. *PLoS ONE* **8**, e79876. (doi:10.1371/journal.pone.0079876)
- Dadvar M, Majd K, Oikonomou E, Fainekos G, Srivastava S. 2022 Joint communication and motion planning for cobots. In *Proc. - IEEE Int. Conf. on Robotics and Automation*, pp. 4771–4777. Philadelphia, PA, USA 23–27

- May 2022. (doi:10.1109/ICRA46639.2022.9812261).
31. Bratman ME. 1987 *Intention, plans, and practical reason*. Cambridge, UK: Harvard University Press.
 32. Baker C, Saxe R, Tenenbaum J, Baker CL. 2011 Bayesian theory of mind: modeling joint belief-desire attribution publication date Bayesian theory of mind: modeling joint belief-desire attribution. *Proc. Annual Meeting of the Cognitive Science Society*, no. 33. Boston, MA, USA. 20–23 July 2011. Cognitive Science Society.
 33. Siebinga O, Zgonnikov A, Abbink DA. 2023 Modelling communication-enabled traffic interactions. Figshare. (doi:10.6084/m9.figshare.c.6632547)
 34. Daamen W, Loot M, Hoogendoorn SP. 2010 Empirical analysis of merging behavior at freeway on-ramp. *Transp. Res. Rec.* **2188**, 108–118. (doi:10.3141/2188-12)
 35. Naumann M, Sun L, Zhan W, Tomizuka M. 2020 Analyzing the suitability of cost functions for explaining and imitating human driving behavior based on inverse reinforcement learning. In *2020 IEEE Int. Conf. on Robotics and Automation (ICRA)*, pp. 5481–5487. IEEE Paris, France. 31 May–31 August 2020. (doi:10.1109/ICRA40945.2020.9196795)
 36. 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. Rob.* **42**, 1405–1426. (doi:10.1007/s10514-018-9746-1)
 37. Xu Y, Bao S, Pradhan AK. 2021 Modeling drivers' reaction when being tailgated: a random forests method. *J. Safety Res.* **78**, 28–35. (doi:10.1016/j.jsr.2021.05.004)
 38. Kondoh T, Yamamura T, Kitazaki S, Kuge N, Boer ER. 2008 Identification of visual cues and quantification of drivers' perception of proximity risk to the lead vehicle in car-following situations. *J. Mech. Syst. Transp. Logist.* **1**, 170–180. (doi:10.1299/jmtl.1.170)
 39. Piao J, McDonald M. 2003 Low speed car following behaviour from floating vehicle data. *IEEE Intelligent Vehicles Symp., Proc.*, pp. 462–467. Columbus, OH, USA. 09–11 June 2003. IEEE. (doi:10.1109/IVS.2003.1212955)
 40. Ji A, Levinson D. 2020 A review of game theory models of lane changing. *Transportmetrica A: Transp. Sci.* **16**, 1628–1647. (doi:10.1080/23249935.2020.1770368)
 41. Li N, Dwyer DW, Zhang M, Yildiz Y, Kolmanovsky I, Girard AR. 2018 Game theoretic modeling of driver and vehicle interactions for verification and validation of autonomous vehicle control systems. *IEEE Trans. Control Syst. Technol.* **26**, 1782–1797. (doi:10.1109/TCST.2017.2723574)
 42. Schwarting W, Pierson A, Alonso-Mora J, Karaman S, Rus D. 2019 Social behavior for autonomous vehicles. *Proc. Natl Acad. Sci. USA* **116**, 24 972–24 978. (doi:10.1073/pnas.1820676116)
 43. de Silva L, Meneguzzi F, Logan B. 2020 BDI agent architectures: a survey. *IJCAI Int. Joint Conf. on Artificial Intelligence*, vol. 2021, pp. 4914–4921. Yokohama, Japan. 7–15 January 2021. ACM.
 44. Premack D, Woodruff G. 1978 Premack and Woodruff : chimpanzee theory of mind. *Behav. Brain Sci.* **4**, 515–526. (doi:10.1017/S0140525X00076512)
 45. Scassellati B. 2002 Theory of mind for a humanoid robot. *Auton. Rob.* **12**, 13–24. (doi:10.1023/A:1013298507114)
 46. Devin S, Alami R. 2016 An implemented theory of mind to improve human-robot shared plans execution. *ACM/IEEE Int. Conf. on Human-Robot Interaction*, vol. 2016, pp. 319–326. Christchurch, New Zealand. 07–10 March 2016. IEEE. (doi:10.1109/HRI.2016.7451768)
 47. Nishi T, Doshi P, Prokhorov D. 2019 Merging in congested freeway traffic using multipolicy decision making and passive actor-critic learning. *IEEE Trans. Intell. Vehicles* **4**, 287–297. (doi:10.1109/TIV.2019.2904417)
 48. Yu H, Tseng HE, Langari R. 2018 A human-like game theory-based controller for automatic lane changing. *Transp. Res. Part C: Emerg. Technol.* **88**, 140–158. (doi:10.1016/j.trc.2018.01.016)
 49. Markkula G, Lin Y-s., Srinivasan AR, Billington J, Kalantari AH, Yang Y, Lee YM, Madigan R, Merat N. 2022 Explaining human interactions on the road requires large-scale integration of psychological theory. *PsyArXiv*. (doi:10.31234/osf.io/hdxbs)
 50. Wan X, Jin PJ, Yang F, Zhang J, Ran B. 2014 Modeling vehicle interactions during merge in congested weaving section of freeway Ramp. *Transp. Res. Rec.: J. Transp. Res. Board* **2421**, 82–92. (doi:10.3141/2421-10)
 51. Kurtc V. 2020 Studying car-following dynamics on the basis of the HighD dataset. *Transp. Res. Record* **2674**, 813–822. (doi:10.1177/0361198120925063)
 52. Siebinga O, Zgonnikov A, Abbink D. 2022 A human factors approach to validating driver models for interaction-aware automated vehicles. *ACM Trans. Human-Rob. Interact.* **11**, 1–21. (doi:10.1145/3538705)
 53. Siebinga Olger. 2023 Software underlying the publication: Modelling communication-enabled traffic interactions (Version 1) [Computer software]. 4TU.ResearchData. (<https://doi.org/10.4121/21196357.V1>)