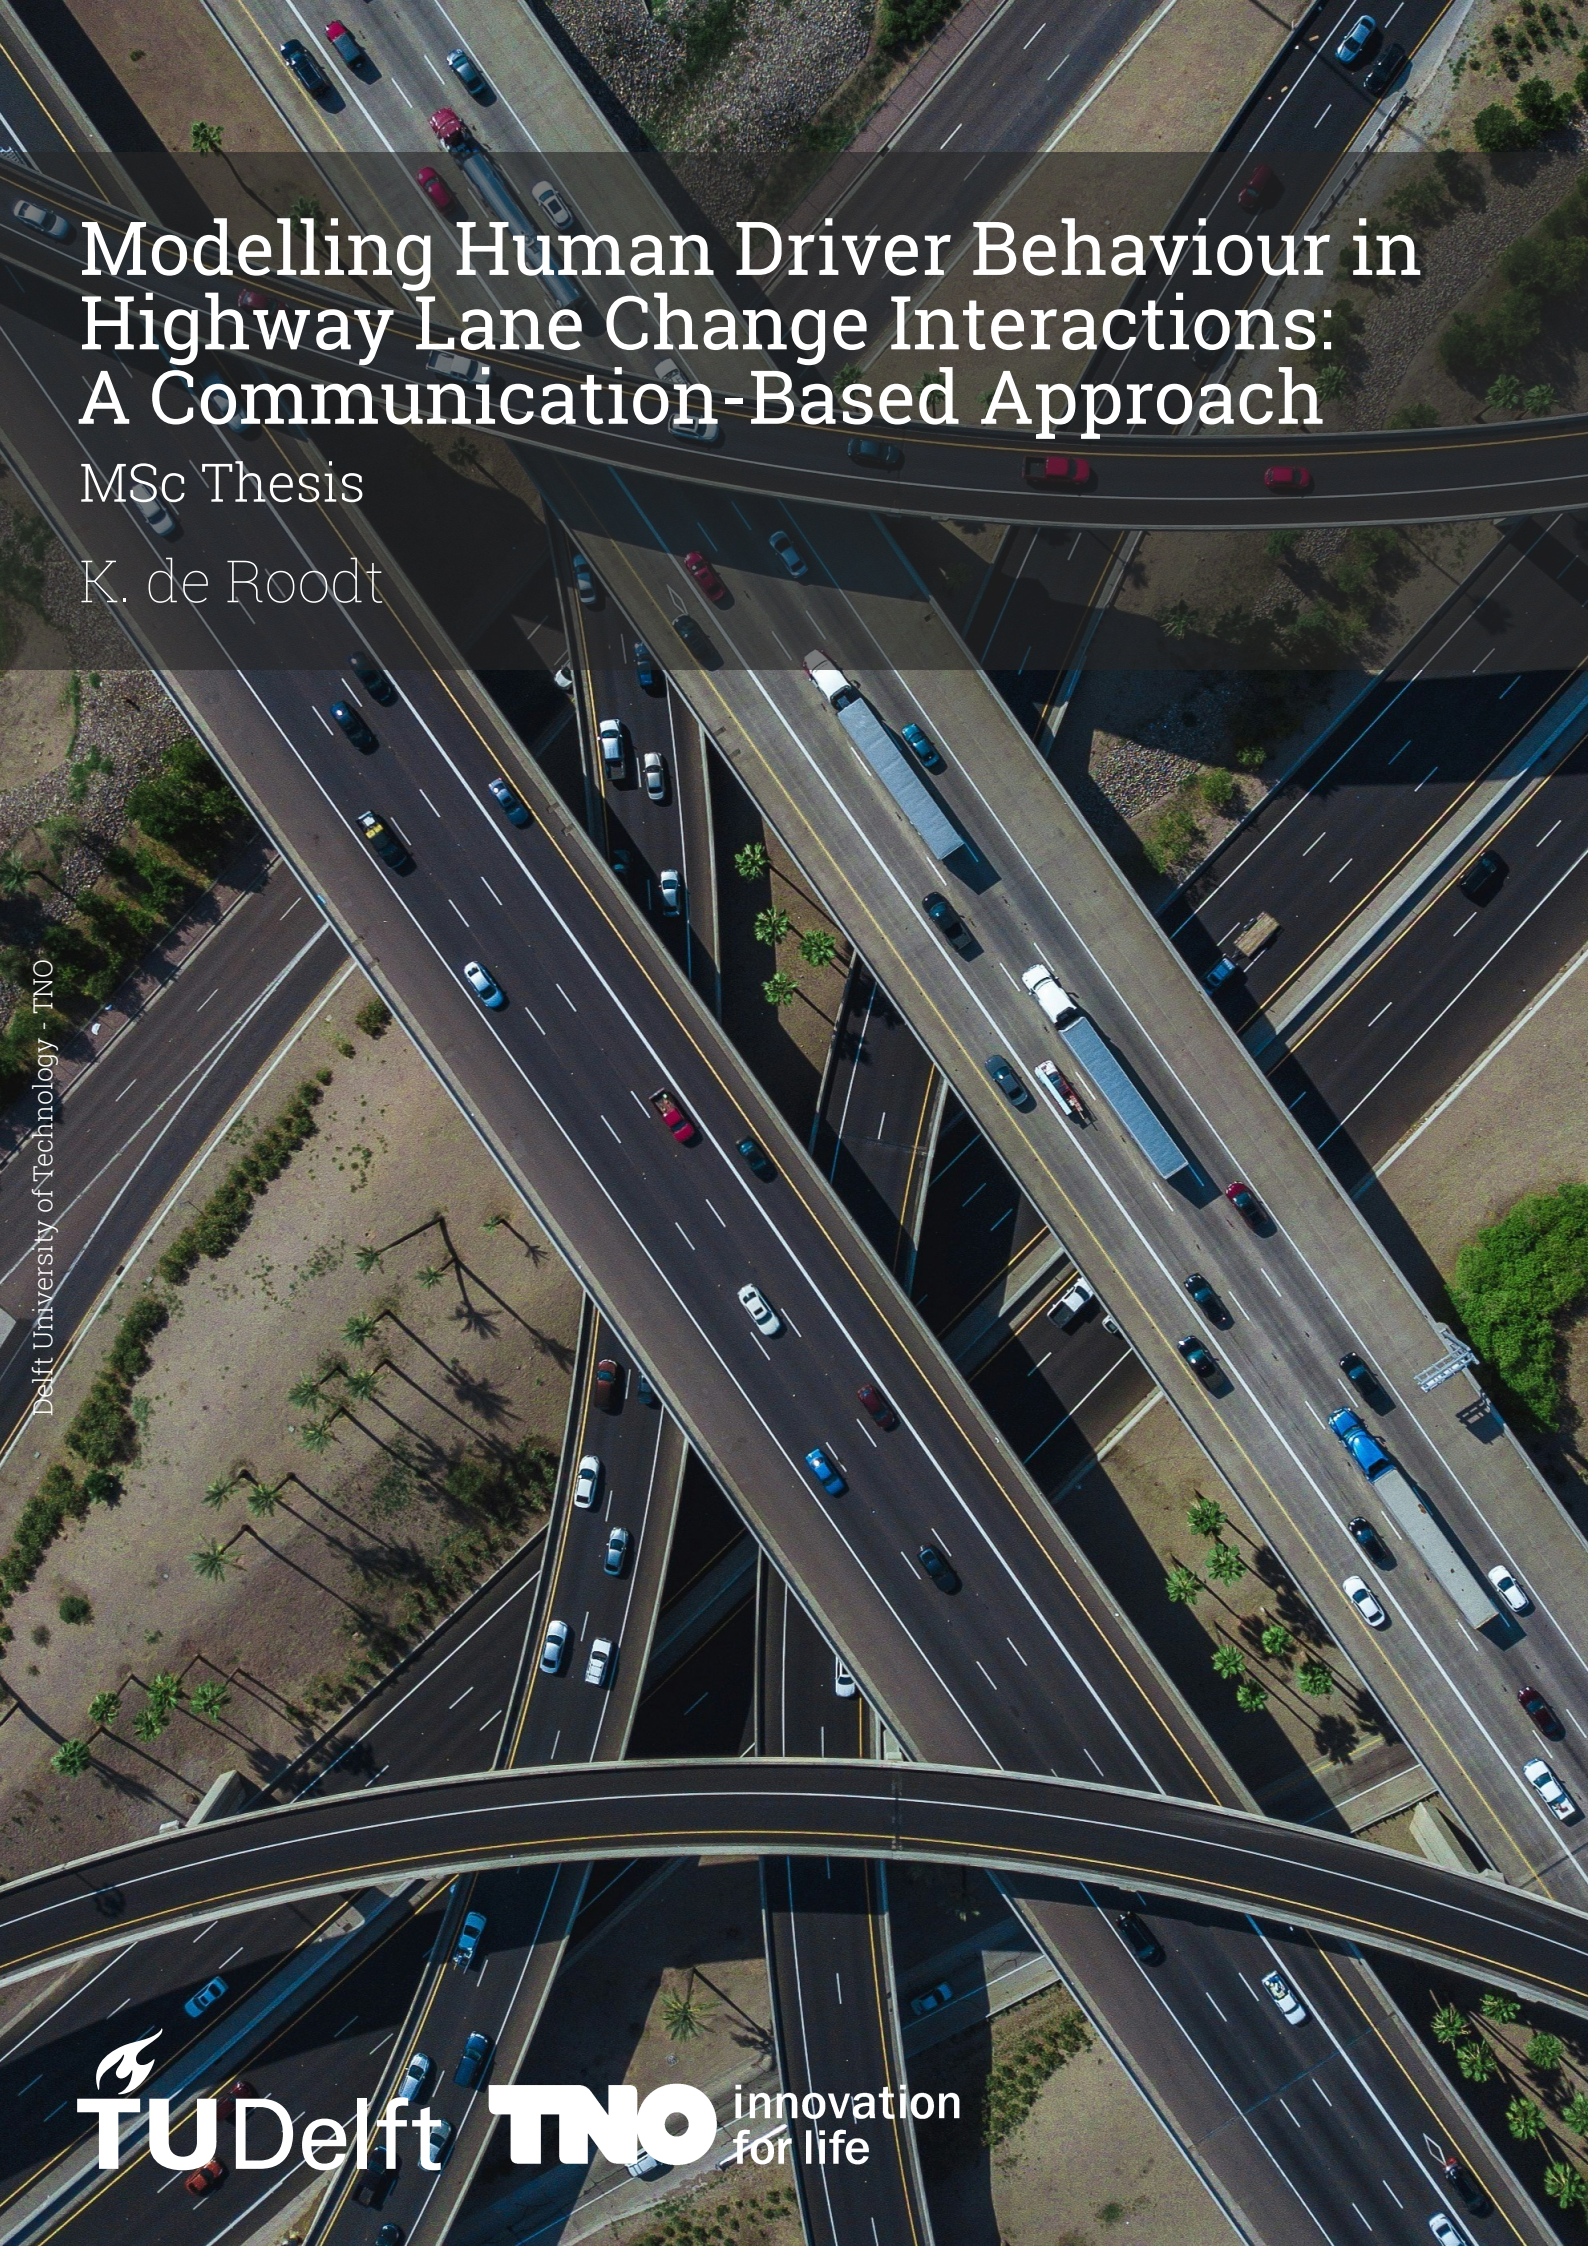

Modelling Human Driver Behaviour in Highway Lane Change Interactions: A Communication-Based Approach

MSc Thesis

K. de Roodt

Delft University of Technology - TNO

TU Delft TNO innovation for life

Modelling Human Driver Behaviour in Highway Lane Change Interactions: A Communication-Based Approach

MSc Thesis

by

K. de Roodt

to obtain the degree of Master of Science at the Delft University of Technology,
to be defended publicly on Friday, February 21, 2025, at 13:30.

Student number:	5520126	
Project duration:	May, 2024 - February, 2025	
Graduation committee:	Dr. A. Zgonnikov,	TU Delft, supervisor
	Dr.Ir. O. Siebinga,	TU Delft, supervisor
	Ir. M. van Weperen,	TNO, external supervisor
	Prof.dr.ir. R. Happee,	TU Delft, external member
	Dr.Ir. I. Martínez,	TU Delft, external member
Faculty:	Faculty of Mechanical Engineering, Delft University of Technology	
Company:	TNO, Department of Integrated Vehicle Safety	
Cover Image:	Aerial Concrete Road by Pexels (modified)	

TU Delft TNO innovation for life

Preface

Changing lanes on a highway might seem like second nature to most humans: you flip on your indicator, glance in the mirror, and seamlessly merge into the adjacent lane—right? However, once I decided to model this “simple” manoeuvre, I quickly realized it was anything but. Over the past year, I have explored lane change theories and driver interaction models, uncovering the complexity behind what appears to be an effortless decision. This work results in a computational lane change model integrating high-level decision-making with low-level control actions into a unified model based on communication and beliefs about other drivers’ intentions. Its validity is assessed against naturalistic driving data, with all findings detailed in this thesis.

Of course, I did not get here on my own. My sincerest thanks go to my supervisors, Arkady, Marijke, and Olger, for their critical and clear guidance on the technical aspects of this project. I am especially grateful to Arkady and Marijke for always showing compassion and mentorship when challenges piled up. With your guidance, I managed to merge through the traffic and stay on track. I also deeply appreciate Olger for his insightful discussions and brainstorming sessions, which played an important role in shaping the ideas in this work.

Last but by no means least, I want to thank my family, friends, and especially my girlfriend, whose unwavering encouragement and constant belief in me have kept me moving forward. This thesis is a reflection of the collective support, advice, and encouragement from everyone who believed in me. I am truly thankful for each and every one of you who helped me reach this milestone.

K. de Roodt
Delft, February 2025

Contents

1 Paper **1**

A Model and Simulation Parameters **21**

B Preliminary Controlled Simulations **23**
B.1 Two Agents: Lead Vehicle Interactions 23
B.2 Three Agents: Left-Following Vehicle Interactions 28
B.3 Four Agents: Complex Multi-Vehicle Interactions 33

C Systematic Validation Against Literature on Lane Change Behaviour **37**
C.1 Effect of Varying Lead Vehicle Gap and Speed 38
C.2 Effect of Varying Left-Following Vehicle Gap and Speed 40

D Grid Search **42**
D.1 Scenario and Dataset Extraction 42
D.2 Grid Search Setup and Parameter Ranges 44
D.3 Performance Metrics 44
D.4 Results 44

E Validation Without Benefit-Based Trigger **45**

1

Paper

Modelling Human Driver Behaviour in Highway Lane Change Interactions: A Communication-Based Approach

K. de Roodt
Delft University of Technology
Delft, The Netherlands

***Abstract*—With the increasing integration of Automated Vehicles (AVs) into our daily traffic, validating their performance poses a significant challenge. Virtual testing, where simulated AVs operate in a simulated environment, has become a widely adopted approach for efficient and cost-effective validation. However, the lack of realistic human behaviour models hinders the realism of these simulations, particularly in simulating reciprocal driver interactions. In this research, I introduce a discretionary lane change model based on the Communication-Enabled Interaction (CEI) framework, which simulates reciprocal driver interactions through implicit communication and belief modelling. These reciprocal interactions involve mutual behaviours, where individual drivers contribute through high-level decisions and low-level control actions. By employing the CEI framework, decision-making and control actions are integrated into a unified model. The proposed model is validated against naturalistic driving data in discretionary lane change scenarios to assess its validity. Results demonstrate that the model successfully reproduces qualitative and quantitative characteristics of human driving behaviour, reflecting both individual behaviours and the collective contributions of multiple drivers. Moreover, it reflects how varied tactical decisions yield distinct, human-like operational execution characteristics. Thereby improving the realism of interactive traffic simulations and posing a step towards improving virtual testing environments for AVs.**

1. Introduction

As Automated Vehicles (AVs) merge into our daily traffic, they promise a future of more efficient commutes, fewer accidents, and an overall improvement in travel quality [1, 2]. Despite rapid advancements, AVs still face challenges interacting with other traffic participants (e.g., pedestrians, passenger cars, or trucks) and incorporating their behaviours into the decision-making process [3]. This issue is particularly evident in scenarios involving *reciprocal* interactions, such as highway lane changes, where road users mutually influence and respond to each other's actions.

In this work, I focus on *discretionary* lane changes, a subset of lane changes performed voluntarily to gain better driving conditions (e.g., a speed advantage) [4]. Compared to mandatory lane changes (e.g., entering a highway), discretionary lane changes are more complex as drivers must evaluate the necessity, desirability, and safety of the manoeuvre [5]. However, drivers only need to evaluate the safety of mandatory lane changes. This makes discretionary lane changes highly sensitive to individual differences and reciprocal influences.

To ensure AVs can manage these interactions safely, leveraging human driver models has emerged as a tool for validating AV performance [6]. Driver models can serve two primary purposes in validating AVs: (1) by acting as performance benchmarks representing competent human drivers (e.g., [7–9]), and (2) by simulating the behaviours of surrounding vehicles to test how an AV responds in various scenarios (e.g., [10–13]).

Most existing performance benchmark models focus on rear-end conflicts in car-following scenarios such as cut-ins, cut-outs, and decelerations (e.g., [14, 15]). However, recent work has begun to address mandatory lane changes. One study proposed a kinematic Bi-Dimensional driver model for evaluating the low-level control of evasive lane changes [8]. Another included decision-making but treated interaction mainly as a one-way response to other vehicles, thus overlooking reciprocal influences [9]. So far, these benchmark models have mainly emphasized single-driver behaviour. They often assume that while the modelled driver reacts to surrounding traffic, those vehicles do not respond in turn. Although this assumption works for car-following models, it does not suit interactive driving scenarios like lane changes. Furthermore, current benchmark models remain limited to mandatory lane changes, leaving discretionary lane changes largely unaddressed.

Besides models specifically meant for benchmarking, there are other more general lane change models (e.g., for modelling surrounding traffic in AV simulation) spanning data-driven or mechanistic modelling approaches. On the data-driven side, large-scale naturalistic datasets and machine learning techniques enable models to learn driver behaviour characteristics (see [16] for a comparison). For instance, the Waymo SimAgents challenge has promoted the development of such data-driven approaches for generating realistic surrounding traffic simulations [17]. Although these models often claim to incorporate interaction, they typically rely on learned correlations (i.e., conditioning one agent's motion on others' behaviours) rather than explicitly modelling reciprocal interactions among drivers. As a result, while they can match observed traffic patterns accurately, they often lack interpretability, making it unclear *why* drivers behave as they do.

Many longstanding mechanistic lane change models exist (e.g., utility-based, game-theory-based, rule-based, etc. See [18] for a review), but they have predominantly operated at the level of traffic microsimulation, focusing on macroscopic traffic flow characteristics rather than individual behaviours. These lane change models often segment the problem into high-level decision-making (i.e., the need to change lanes), followed by gap acceptance (i.e., evaluating whether the available space is sufficient), and end with low-level control actions (e.g., changes in velocity). Research on a driver's motivations for discretionary lane changes has focused on variables like relative distance and speed [12], acceleration [13], or including variables such as average travel time and time gaps [19]. To model gap acceptance and control actions, these approaches often utilize gap acceptance theory (e.g., [5, 12, 13]) alongside acceleration models (e.g., the intelligent driver model [20]) to simulate behaviour before and after lane changes. However, these models primarily focus on decision-making aspects, overlooking the execution phase of the lane change, particularly for discretionary lane changes [21]. Despite the significant impact that the execution of lane changes has on driver behaviour in real-world traffic, researchers rarely capture the low-level control actions [18]. As a result, this limits their ability to reproduce the variability (the range of human behaviours that can be expected in response to a given situation [22]) arising from different tactical behaviours (i.e., describing *which* manoeuvres are executed [23]) in the resulting operational behaviour (i.e., describing *how* manoeuvres are executed) within a unified framework. Furthermore, they often ignore or simplify interactions into one-way responses, overlooking the reciprocal influences [18]. Therefore, it is challenging for traditional, single-agent-focused lane change models to reproduce human interactive behaviour accurately.

In an attempt to model drivers as a joint interactive system (as opposed to individual drivers), researchers have widely used game-theoretic approaches (e.g., [10, 11], see [24] for a review). However, these models often assume humans act as rational, utility-maximizing agents without any form of communication. Furthermore, most game-theory-based interaction models rely on discrete actions for drivers. This approach is useful for higher-level decision-making (e.g., yielding or lane changing) but fails to encompass the low-level control inputs (e.g., changes in velocity). As noted by Siebinga *et al.* [25], these assumptions do not accurately reflect real-world driver behaviour.

Based on these limitations, there is a need for an interpretable, unified model that can describe reciprocal interactions and integrate decision-making and low-level control actions. I aim to address this gap by developing a computational human driver model based on the Communication-Enabled Interaction (CEI) framework [25], see Section 2 for a brief overview. This framework models the two-way interaction between drivers by considering the joint interactive system, guided by communication and beliefs about other drivers' intentions.

The CEI framework has recently been validated only for a simple one-dimensional merging scenario without lateral behaviour [26]. Therefore, I extended the CEI framework to more realistic scenarios involving two-dimensional dynamics, multiple agents, and multiple lanes, as detailed in Section 3. I validated whether and under which circumstances the model could reproduce individual and mutual behaviours at both tactical and operational levels, as described in Section 4. To do so, I compared the driving characteristics from naturalistic driving data with the behaviour generated by the model in a simulation environment. Combined, the contributions of this paper are threefold: (1) the development of a computational model for discretionary lane changes based on the CEI framework, (2) the validation of the model against naturalistic driving data, and (3) the demonstration of the model to reproduce individual behaviours and collective contributions at both tactical and operational levels within a unified model.

2. Background: Communication-Enabled Interactions Framework

Siebinga *et al.* [25] introduced the CEI framework to model reciprocal interactions by considering how drivers communicate and adjust their behaviours based on perceived risks and probabilistic beliefs about others. The framework comprises four components: deterministic plan, communication, probabilistic belief, and risk perception, as illustrated in Figure 1.

2.1. Plan

Each driver formulates a deterministic plan that outlines its intended actions over a future time horizon. This plan uses a cost-minimization approach that prioritizes objectives such as maintaining a desired speed and minimizing acceleration control inputs. During this optimization, the plan is subject to a risk constraint to reduce the perceived risk with the updated plan.

2.2. Communication

Drivers communicate their plans or intentions either explicitly (e.g., turn indicators) or implicitly (e.g., position, velocity, or acceleration). Other drivers interpret these signals to form beliefs about the communicator's intended actions. In this way, communication links one driver's plan to the other driver's belief.

2.3. Belief

Drivers use the observed cues to generate probabilistic beliefs about one another's positions in the near future, which are modelled as normal Gaussian distributions. At every time step, each driver's belief is represented as a probability density over the other drivers' future positions. By combining these beliefs with their deterministic plans, drivers compute the probability of collisions, leading to a perceived risk.

2.4. Risk-Based Replanning

The CEI framework integrates risk-based replanning as a form of *bounded rationality* and *satisficing* to reflect

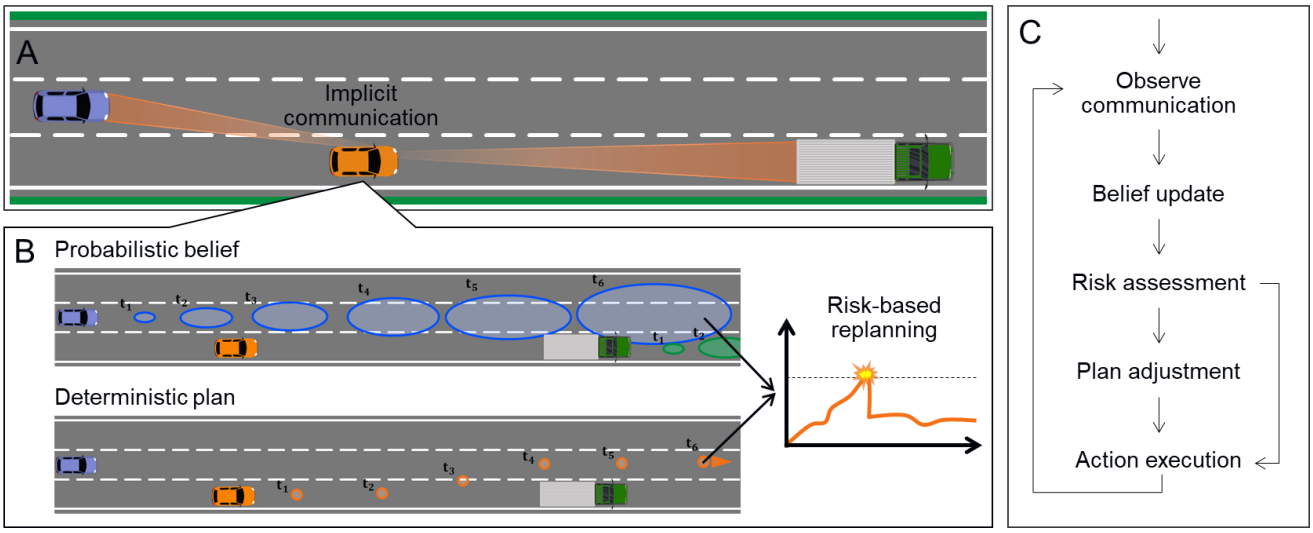

Figure 1: An overview of the CEI-based lane change model. **A)** A typical discretionary lane change scenario of the ego vehicle (orange) overtaking the slower lead truck (green) while evaluating the observable cues of the left-following vehicle (blue). **B)** Each driver is modelled using the four CEI components: plan, communication, belief, and risk. Drivers have deterministic plans for their own actions and hold probabilistic beliefs about others' future positions based on implicit communication. Combined, plan and belief results in a risk perception. Replanning is triggered when the risk threshold is exceeded based on the risk perception or when the adjacent left lane is more beneficial, allowing drivers to initiate lane changes proactively before risks become critical. While only the orange vehicle's components are illustrated, the framework applies to all vehicles; the blue vehicle and truck have the same components. **C)**: The sequence of operations drivers perform each timestep, where plan adjustment is only performed if the perceived risk exceeds the threshold or when the adjacent left lane is more beneficial.

drivers' cognitive and perceptual characteristics. Drivers continuously assess the risk associated with their current plans by estimating potential hazards, such as proximity to other drivers. Each driver has a personal risk threshold based on subjective evaluations. When the perceived risk exceeds this threshold, the driver initiates a replanning process to adjust its trajectory and reduce the risk to an acceptable level.

Using a form of satisficing, drivers seek an optimal solution only when the perceived risk exceeds their personal threshold and follow the plan while the risk remains below it. This reduces cognitive load by avoiding constant optimization. Drivers with lower risk thresholds tend to adjust their plans earlier, while those with higher thresholds may delay adjustments, potentially relying on others to mitigate the risk first. This reflects how drivers balance safety and efficiency within their cognitive limits.

2.5. Sequence of Operations

By integrating all four framework components, drivers perform the following sequence of operations at each timestep:

- Observe communication: drivers observe the current state of others.
- Belief update: drivers update their probabilistic beliefs based on the observed communication.
- Risk assessment: drivers estimate the risk associated with their current plans using the updated beliefs.
- Plan adjustment: drivers replan their trajectory to mitigate the risk only if the perceived risk exceeds the upper-risk threshold.
- Action execution: drivers execute the next step of their plan, and the process repeats in the subsequent time interval.

3. Communication-Enabled Interactions Model

3.1. Model Development

The proposed two-dimensional lane change model is based on the CEI modelling framework [25] and its subsequent implementation in [26]. The model consists of instantiations for each component, as shown in Figure 1. This involves the overall dynamics, planning, communication, belief construction, risk evaluation, and an additional replanning mechanism for early lane changes.

3.1.1. Overall Dynamics. The overall dynamics of the simulation are extended to two dimensions, as opposed to the original one-dimensional point mass kinematics in [26]. The agents drive in the 2-dimensional x-y-plane and are modelled using a dynamic bicycle model [27]. The dynamic bicycle model is constructed using six states and two inputs:

$$X = \begin{bmatrix} x & y & v_{long} & v_{lat} & \psi & \dot{\psi} \end{bmatrix}^\top \tag{1}$$

$$U = \begin{bmatrix} a_{long} & \delta_f \end{bmatrix}^\top \tag{2}$$

where x and y denote the longitudinal and lateral position, respectively, v_{long} and v_{lat} are the corresponding velocities, and ψ denotes the yaw angle. The inputs of the system are the longitudinal acceleration a_{long} and the front wheel steering angle δ_f, with maximum absolute values of $a_{long} = 1.5\ m/s^2$, $\delta_f = \frac{\pi}{18}\ rad$.

3.1.2. Plan. The planning is performed in acceleration space and consists of a set of waypoints over a time horizon $T = 8.0s$. The plan is obtained by optimizing cost function J_{plan} for each waypoint of the time horizon:

$$J_{plan} = \sum_{t=t_0}^{T} \left(\lambda_1 (dv_{long})^2 + \lambda_2 (dv_{lat})^2 + \lambda_3 (d\psi)^2 + \lambda_4 \delta_t^2 + \lambda_5 a_t^2\right) \tag{3}$$

where dv_{long} and dv_{lat} are the deviation from the desired longitudinal and lateral velocity, $d\psi$ is the deviation from the desired heading, δ_t is the steering input, and a_t is the acceleration input. The λ-parameters are weights to adjust the relative importance of the different components. By adjusting the cost-function weights, the model can be tuned to favour lane changes over speed reductions (or vice versa), achieving different tactical and operational behaviours.

The cost function does not contain terms for collision avoidance or maintaining within lane bounds. Instead, the CEI framework ensures safety by requiring drivers to keep the risk below the risk threshold, which is enforced through an optimization constraint in the model's planning process.

At each simulation timestep, the first action (acceleration and steering) from the updated plan is applied to the bicycle model and removed from the planning horizon. Because drivers do not continuously replan every timestep, I extend the acceleration input using the final action to reflect the idea of *continuing with the current plan*. For steering, I compute a new steering input using a PD controller to maintain lane alignment and counteract lateral drift, which can arise from minimal steering inputs at high speeds near the end of the planning horizon.

3.1.3. Communication. Drivers base their belief about other agents' future trajectories on implicit communication through vehicle movements. At each time step, drivers observe the positions, velocities, and accelerations of surrounding vehicles. While position and acceleration are assumed to be observed as perfect, velocity perception is modelled as noisy and influenced by perceptual delays. This accounts for drivers' tendencies to under- or overestimate other vehicles' velocities at high speeds [28, 29]. The updates are defined as:

$$\Delta v_{\text{long},o} = \alpha_c \left(v_{\text{long},t} - v_{\text{long},o}\right) + \beta dW \tag{4}$$

$$\Delta v_{\text{lat},o} = \alpha_c \left(v_{\text{lat},t} - v_{\text{lat},o}\right) + \beta dW \tag{5}$$

The perception update rate α_c controls how quickly drivers adjust their perceived velocities toward the true values (subscript t), modelling the perception delay. The noise scaling factors (β_{long} and β_{lat}) introduce variability into the perception, simulating the inaccuracy in estimating other vehicles' velocities. Additionally, W represents a stochastic Wiener process, further contributing to the noise in perception. This method is adapted from [26].

3.1.4. Belief. The observed communication is used to create a two-dimensional belief about the future positions of other drivers. These beliefs consist of belief points, which are probabilistic predictions of the other vehicle's position at specific future times. The belief points share the same time horizon and planning frequency ($4Hz$) as the driver's own plan. Each belief point is constructed using a joint longitudinal and lateral probability distribution, assuming independence between the longitudinal and lateral components.

The longitudinal belief is based on the CEI simple merging model [26]. Each of the longitudinal belief points is represented by the sum of two normal distributions:

$$b_t = \frac{1}{2}\mathcal{N}\left(\mu_t, \sigma_t^2\right) + \frac{1}{2}\mathcal{N}\left(\mu_t, \phi\sigma_t^2\right) \tag{6}$$

where b_t is the belief point representing the probability distribution over longitudinal positions for the other vehicle at time t. ϕ is a scaling factor that adjusts the variance of the second normal distribution to capture uncertainty in acceleration behaviour.

The expected position μ_t and variance σ_t^2 for each belief point are calculated by extrapolating the other vehicle's observed position, velocity, and heading:

$$\mu_{x,t_b} = x_{t_0} + (v_{long,t_0}\cos(\psi_{t_0}) - v_{lat,t_0}\sin(\psi_{t_0}))\, t_b + \frac{1}{2}(t_b - t_0)^2 \mu_{x,a} \tag{7}$$

$$\sigma_{x,t_b}^2 = \left(\frac{1}{2}\left(t_b - t_0\right)^2\right)^2 \sigma_{x,a}^2 \tag{8}$$

where μ_{x,t_b} denotes the extrapolated x position at time t_b, x_{t_0} denotes the observed position, v_{long,t_0} and v_{lat,t_0} the perceived (noisy) velocity of the other vehicle, t_b denotes the time of this belief point and t_0 the current time. The parameters $\mu_{x,a}$ and $\sigma_{x,a}^2$ are derived from a normally distributed expected acceleration based on the driver's memory of recent longitudinal acceleration observations, as detailed in [26].

The lateral belief is modelled using the same method as the longitudinal belief but includes an adjustment for the vehicle's heading to prevent excessive lateral drift. This adjustment converges the belief toward lane centring by the end of the planning horizon. The adjusted heading ψ_{conv} is computed as:

$$\kappa_{conv} = \frac{1}{1 + \exp(-1(\frac{t_b}{T} - \eta_{conv}))} \tag{9}$$

$$\psi_{\text{conv}} = (1 - \kappa_{conv})\psi \tag{10}$$

where η_{conv} denote the shift parameter where κ_{conv} starts to adjust the heading.

This approach avoids extrapolating lateral positions into unrealistic trajectories. The expected lateral position μ_{y,t_b} and variance σ_{y,t_b}^2 for each belief point are calculated by extrapolating the other vehicle's observed position, velocity, and adjusted heading:

$$\mu_{y,t_b} = y_{t_0} + (v_{long,t_0} \sin(\psi_{conv,t_0}) + v_{lat,t_0} \cos(\psi_{conv,t_0}))\, t_b + \frac{1}{2}(t_b - t_0)^2 \mu_{y,a} \tag{11}$$

$$\sigma^2_{y,t_b} = \left(\frac{1}{2}(t_b - t_0)^2\right)^2 \sigma^2_{y,a} \tag{12}$$

where μ_{y,t_b} denotes the extrapolated y position at time t_b, y_{t_0} denotes the observed lateral position, and $\mu_{y,a}$ and $\sigma^2_{y,a}$ are derived from a normally distributed expected lateral acceleration based on the driver's memory of recent lateral acceleration observations.

3.1.5. Risk. Risk perception involves the driver's planned trajectory, beliefs about other vehicles' future positions, and maintaining safe lateral positions within lane boundaries. Combined, three elements contribute to risk perception: collision risk (P_{col}), driving beyond the road boundaries (P_{bounds}), and lane proximity risk (P_{prox}).

At each time step, every corresponding plan-belief point pair is evaluated against the three risk components, resulting in a perceived risk contribution for each component ranging between 0 and 1. The combined risk value consists of the summation of all three risk elements.

Collision Risk: The collision risk (P_{col}) is derived from the planned position and the belief point representing the perceived probability of vehicle positions at the same time step. This method is adapted from [26]. To estimate the probability of a collision, the concept of *bounds of collision* is used [25]. These bounds define the positions of the other vehicle that would lead to a collision based on the position of the ego vehicle. At each belief time point t_b, the collision risk is determined by the probability that the other vehicle's position falls within the collision bounds of the ego vehicle:

$$P_{col}\left(x_{e,t_b} - r_l \leq x_{b,t_b} \leq x_{e,t_b} + r_l\right) \cdot P_{col}\left(y_{e,t_b} - r_w \leq y_{b,t_b} \leq y_{e,t_b} + r_w\right) \tag{13}$$

where x_{b,t_b} and y_{b,t_b} are the belief x and y-positions of the other vehicle at time t_b, x_{e,t_b} and y_{e,t_b} are the plan x and y-position of the ego vehicle, added and subtracted by the vehicle length r_l and vehicle width r_w to obtain the upper and lower bounds of collision.

Road Boundary Risk: The risk of driving beyond the road boundaries (P_{bounds}) increases as the lateral position of the ego vehicle approaches the sides of the road. This risk is modelled using hyperbolic tangent functions, which smoothly transition from low risk when the vehicle has a large distance to the road edge to high risk as it comes close to the road edge, based on the planned point's lateral position x_{t_b}. The road boundary risk is calculated as:

$$P_{\text{bounds}} = 1 - \frac{1}{2}\tanh\left(\eta(x_{e,t_b} + \frac{r_w}{2})\right) + \frac{1}{2}\tanh\left(\eta(x_{e,t_b} - \frac{r_w}{2})\right) \tag{14}$$

where η is the scaling factor that determines the steepness of the function. r_w represents the vehicle's width to let the risk be approximately one as the vehicle approaches the road edge and decreases toward 0 when the vehicle is centred within the lane.

Lane Proximity Risk: The lane proximity risk (P_{prox}) evaluates the risk of driving too close to adjacent lanes without intending a lane change (e.g., in the case of lane-keeping). This risk is also modelled using hyperbolic tangent functions to estimate proximity to each adjacent lane's centre (x_{centre}) based on the planned point's lateral position x_{t_b} and the vehicle width r_w. For each lateral position x_{t_b}, the positional risk is calculated as:

$$P_{\text{prox, pos}} = \frac{1}{2}\tanh\left(\eta(x_{e,t_b} - x_{\text{centre}} + \frac{r_w}{2})\right) - \frac{1}{2}\tanh\left(\eta(x_{e,t_b} - x_{\text{centre}} - \frac{r_w}{2})\right) \tag{15}$$

An additional Gaussian term, $P_{\text{prox, head}}$, lowers the risk if the vehicle's heading deviates from the road's direction, ensuring that when a lane change is actually intended, the model does not mistakenly discourage it. The addition heading term that lowers P_{prox} is calculated as:

$$P_{\text{prox, head}} = \exp\left(-0.5 \cdot \left(\frac{\psi_{t_b}}{\sigma_{heading}}\right)^2\right) \tag{16}$$

where $\sigma_{heading}$ determines how sensitive the risk is to deviations of the vehicle's heading (ψ_{t_b}).

The combined risk for driving too close to adjacent lanes is the multiplication of the positional and heading risks:

$$P_{\text{prox}} = P_{\text{prox, pos}} \cdot P_{\text{prox, head}} \tag{17}$$

Risk Incentive: The perceived risk is assessed relative to two dynamic thresholds ρ_u (upper threshold) and ρ_l (lower threshold). These dynamic thresholds use an incentive function that reflects traffic conventions. The underlying idea is that while two drivers may perceive the same level of risk, traffic conventions can influence which driver is more likely to take action. For instance, in car-following scenarios, it is typically the responsibility of the following vehicle to avoid a collision. Each threshold consists of a driver's individual base value (θ_u for the upper threshold and θ_l for the lower threshold), adjusted by an incentive function:

$$\rho_u = \theta_u + \lambda_{u,1}\Delta x + \lambda_{u,2}\Delta v_{long} \tag{18}$$

$$\rho_l = \theta_l + \lambda_{l,1}\Delta x + \lambda_{l,2}\Delta v_{long} \tag{19}$$

where Δx and Δv_{long} are the longitudinal position and velocity difference between the ego and a leading and/or lead vehicle. These adjustments apply only to vehicles in the current lane, keeping risk thresholds for other lanes unchanged to maintain equal bounds.

A replan is triggered when the upper-risk threshold (ρ_u) is exceeded, aiming to reduce the perceived risk below $0.8\rho_l$. If the risk stays below the lower risk threshold (ρ_l) for longer than the saturation time (τ), a replan ensures the driver does not remain with an overcautious plan, constraining risk below $0.5\rho_u$. Another replan occurs if the desired velocity is reached to maintain it. This method is adapted from [26].

3.1.6. Benefit-Based Replanning. Human drivers often initiate discretionary lane changes *before* risk levels become critical, particularly when they foresee speed gains in an adjacent lane [19, 30]. To model this behaviour, I introduce an additional *benefit-based* trigger alongside the existing risk-based trigger. This trigger uses a speed incentive [12] combined with evidence accumulation to model response timing, which has been previously utilized in studies on driver decision-making [31, 32]. The benefit-based trigger considers how much faster the adjacent lane is compared to the driver's current lane, accumulating evidence toward a "change lanes" decision (see Figure 2) whenever the speed advantage remains sufficiently high.

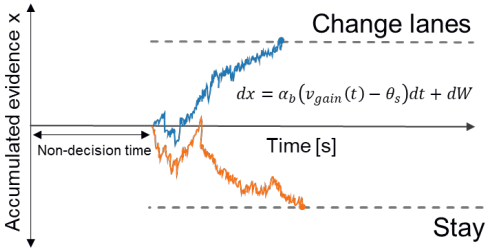

Figure 2: A visual representation of the evidence accumulation model for driver decision-making illustrating how evidence favouring a "change lanes" or "stay" decision accumulates over time. This figure is adapted from [32].

The evidence accumulation consists of variable drift rate $s(t)$ and constant boundary $b(t)$. The drift rate depends on the speed incentive v_{gain} and parameters α, θ_s, and W. This drift rate represents the relative evidence x favouring either a "Change lanes" or "Stay" decision at any given time. The "Stay" decision resets evidence accumulation, preventing evidence from continuously growing negatively. This reset allows the model to reassess lane change decisions when conditions change quickly.

The parameter α_b is the contribution of incoming perceptual information to the accumulated evidence, while θ_s sets the threshold of v_{gain} at which the drift rate changes sign. W denotes a stochastic Wiener process (thus, dW is a sample from a normal distribution $\mathcal{N}(\mu = 0, \sigma^2 = dt)$):

$$s(t) = \alpha_b \left(v_{gain}(t) - \theta_s\right) + dW \tag{20}$$

The drift process ends upon reaching the positive or negative boundary, where height B represents the required amount of evidence to make a decision. In previous research, constant [32, 33] or collapsing boundaries [32, 34] were used. Recent evidence concluded that simpler evidence accumulation models with a constant boundary can accurately account for decision outcomes, similar to more complex models [32, 33]. For this reason, a constant boundary to describe the end of the drift process is used:

$$b(t) = \pm B \tag{21}$$

The benefit of driving in the adjacent lane is estimated using the anticipated speed to assess its advantage:

$$v_{\text{gain}} = v_{\text{ant}}^{i-1} - v_{\text{ant}}^{i} \tag{22}$$

where $i-1$ and i are the left adjacent and current lanes, respectively. The remaining details of the formulation of anticipated speed v_{ant} can be found in [12].

The benefit-based trigger is integrated into the plan as a soft constraint, allowing discretionary manoeuvres when beneficial without enforcing lane changes universally. The constraint is applied with $\gamma = 1$, aligning the final lateral position y_T with the centre of the desired lane y_d. For all other cases, the constraint is deactivated with $\gamma = 0$:

$$J = J_{plan} + \gamma \left(y_T - y_d\right)^2 \tag{23}$$

3.2. Model and Simulation Parameters

The model and simulation rely on 24 parameters. A list of all parameters, including their values and descriptions, is provided in Appendix A. Of these, 14 parameters were manually set based on design considerations, while the remaining 10 were tuned using preliminary simulations and a grid search. The preliminary simulations are used to tune and evaluate model parameters under controlled scenarios, demonstrating the model can reproduce expected behaviours before comparison against naturalistic driving data (details in Appendix B).

The timing parameters for the simulation include the time step (Δt), planning horizon (T), and belief memory duration (T_m). These were designed to suit the scenario, assuming drivers complete a lane change manoeuvre within the planning horizon and to ensure reasonable computation times.

For the communication parameters, the noise scaling factors (β_{long} and β_{lat}) for communication were manually set to reflect a reasonable level of noise. The perception update rate (α_c) was based on an empirical study on driver perception and response times [28].

For the belief parameters, the comfortable accelerations ($a_{c,\text{long}}$ and $a_{c,\text{lat}}$) were derived based on empirical studies [35]. The belief scaling factor (ϕ) was adjusted to ensure adequate resolution in the risk signals, with a chosen value of $\phi = 3.0$.

The cost-function weights (λ_1 to λ_5) were manually adjusted based on a systematic validation of the model using the controlled simulations (Appendix B). Using the cost-function weights, the model balances between different driving objects (e.g., reduce speed or perform lane changes) in the planning process and to obtain stable behaviour.

To account for variability among drivers due to differences in risk tolerance and decision-making styles, I modelled the risk thresholds (θ_u, θ_l) as uniformly distributed across the driver population. The upper threshold θ_u was uniformly distributed between 0.4 and 0.5, and the lower threshold θ_l between 0.2 and 0.3.

The drift rate α_b and parameter θ_s were fitted using grid search to reproduce observed tactical behaviours in a preliminary scenario (details in Appendix D). To account for drivers' variability due to individual decision-making styles, α_b was uniformly distributed between 0.02 and 0.04, and θ_s between 3.0 and 4.0 across the driver population.

4. Validation Against Naturalistic Driving Data

To demonstrate that the proposed driver model reproduces human-like behaviour, I validated it against naturalistic driving data for two representative discretionary lane change scenarios. These scenarios illustrate how individual drivers with unique preferences respond to a slower lead vehicle. By targeting these scenarios, I aim to demonstrate that the proposed approach leads to human-like variability in tactical behaviour (i.e., *which* manoeuvres are executed) and operational behaviour (i.e., *how* the manoeuvres are executed).

While these scenarios represent only a subset of real-world driving situations, they demonstrate the fundamental workings of the approach. Additional preliminary checks using controlled simulations for other scenarios are provided in Appendix B. A systematic validation through controlled simulation against empirical research on lane change behaviour is included in Appendix C.

4.1. Validation Setup

For the validation, I used the highD dataset [36], which contains high-precision vehicle trajectories from German highways captured by unobtrusive aerial drones. Two discretionary lane-change scenarios are selected from the dataset, each with distinct incentives to overtake or continue in the current lane. Both scenarios include at least a slow lead vehicle in the right lane, an approaching ego vehicle, and a following vehicle in the left lane, but they differ in how favourable overtaking is:

- Scenario A high-speed-difference (Figure 3A): A slow-moving lead vehicle (green) occupies the right-most lane, while a faster ego vehicle (orange) approaches from behind. A left-following vehicle (blue) travels in the adjacent lane, and a left-leading vehicle (purple) is present but does not directly influence the lane change. In this scenario, multiple responses are feasible for the ego, such as accelerating and changing lanes, braking to let the left-following vehicle pass before changing lanes or continuing behind the slower lead vehicle.
- Scenario B low-speed-difference (Figure 3B): A similar configuration to scenario A, except the ego vehicle's speed is closer to that of the lead vehicle, reducing the urgency and potential benefit of overtaking.

Human drivers exhibit considerable variation in tactical and operational behaviour even under these seemingly similar conditions [22, 37], e.g., some overtake promptly, while others continue following for extended periods. To capture and validate the variability in behaviours, I employ two complementary methods:

- Hausdorff extraction method [22]: a mathematical metric (Hausdorff distance) that identifies and extracts traffic scenarios from large-scale datasets based on their similarity to a scenario.
- Human factors validation approach [38]: this framework outlines validation steps to assess how a driver model aligns with naturalistic driving data.

4.1.1. Hausdorff Extraction Method. To systematically obtain data reflecting human variability in the two scenarios, I used the Hausdorff extraction method [22]. First, each scenario (A or B) was formalized as a *scenario of interest*, specifying the relative positions and velocities of the ego and surrounding vehicles. These were then converted into a *context set*, which mathematically encodes multi-dimensional data based on vehicle positions and velocities relative to the ego vehicle.

The Hausdorff distance measured the dissimilarity between the scenario of interest and other context sets in the highD dataset. It calculated the greatest minimum distance between points in the two sets, identifying contexts most similar to the scenario of interest. The 300 closest samples were then selected for each scenario.

Figure 3 (C and D) illustrates the spatial and velocity distribution of vehicles in these extracted samples. In Scenario A, 108 out of 300 samples involve a lane change (36%), while 192 correspond to lane-keeping. In Scenario B, 45 samples involve a lane change (15%), and 255 correspond to lane-keeping.

4.1.2. Human Factors Validation Approach. The validation starts with a scenario-level validation of scenarios A and B to illustrate how the model's internal components (plan, communication, risk, and belief) interact to produce tactical decisions and their operational execution.

After the scenario-level validation, I apply a two-stage human factors validation approach outlined in [38]. Using this approach, the model's tactical and operational behaviour is validated across all 300 representative extracted samples per scenario. In the first stage (tactical validation), I validate whether the model's high-level decisions (e.g., lane change vs. lane keeping) align with the variability observed in naturalistic data. In the second stage (operational validation), I examine whether the model's lane-change and lane-keeping behaviours reflect the empirically observed patterns that emerge once a driver commits to overtaking or continues to follow a slower lead vehicle.

4.1.3. Simulation Setup. A simulation environment replicating highD highway conditions was created. The simulation is implemented in a two-dimensional plane, representing a straight highway segment of fixed length ($420m$) and a varying lane width and number of lanes (2 or 3) depending on the original dataset recording.

The initial positions and velocities of all vehicles are based on their Hausdorff-extracted locations and observed speeds in the data sample, respectively. Each agent's preferred velocity is then set to its maximum velocity, defined as the point at which the driver no longer accelerates. Risk thresholds and evidence accumulation parameters are uniformly distributed per sample, while other parameters remain constant. All simulations include the benefit-based trigger. For completeness, Appendix E provides results from the same scenarios without this trigger.

The process of simulating each scenario (A and B), containing 300 samples each, was repeated 10 times to capture the variability in behaviour through the stochastic components of the model and the uniformly distributed parameters. This yielded 6000 simulations in total.

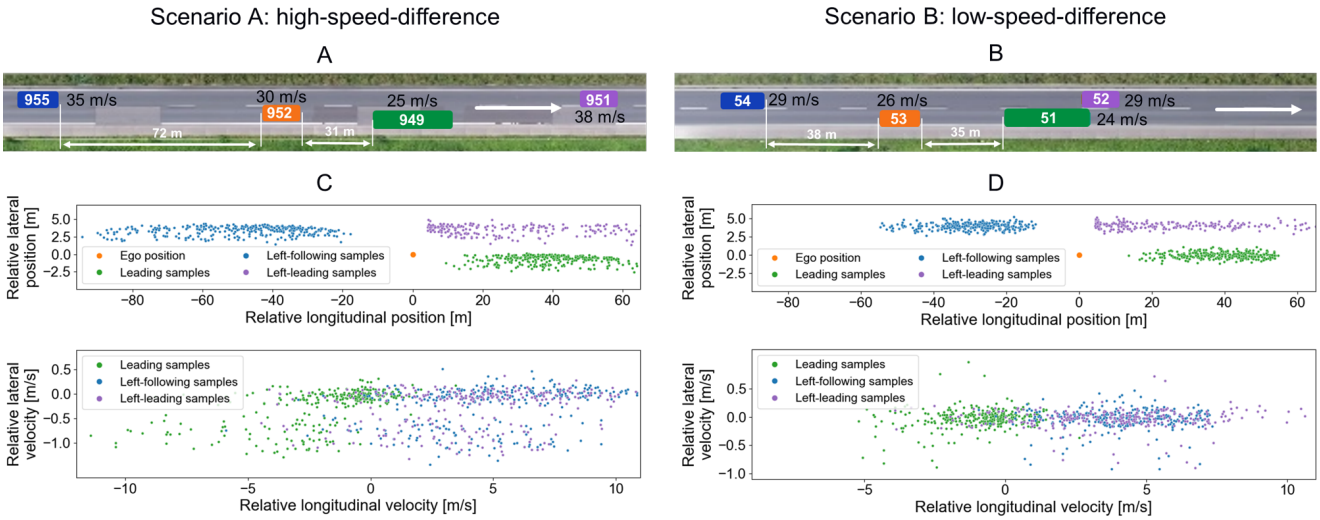

Figure 3: An overview of the scenarios and the spread of representative samples for each scenario. **A-B)** Scenario A (highD dataset 24, frame 17620, ego vehicle id 952) and scenario B (highD dataset 22, frame 1068, ego vehicle id 53). Both scenarios contain an ego (orange), leading (green), left-following (blue), and a left-leading (purple) vehicle. The white arrow denotes the direction of travel. The blue vehicle denotes a left-following vehicle, the green vehicle denotes a leading truck, and the purple vehicle denotes a left-leading vehicle. **C-D)** The spread of the points represents the samples obtained after the Hausdorff extraction [22], for scenario A (first column - C) and scenario B (second column - D). The upper plot displays the positions of surrounding vehicles relative to the ego vehicle, while the lower plot illustrates their relative velocities. The green, blue, and purple dots represent the 300 closest samples extracted from the highD dataset.

4.2. Scenario-Level Validation

This section qualitatively analyzes individual driver behaviours and model dynamics in scenarios A and B. The left-leading vehicle (purple vehicle in Figure 3A-B) is excluded from this comparison as it does not directly influence the lane-change interaction, ensuring a clearer focus on the primary drivers involved. Additional qualitative results on individual behaviours in controlled simulations are provided in Appendix B.

4.2.1. Scenario A: High-Speed-Difference. In scenario A (Figure 4A), three vehicles travel on a two-lane highway. The ego vehicle (orange) initially approaches the lead vehicle (green) with a relative speed of approximately $5m/s$. At the same time, the left-following vehicle (blue) has a similar $5m/s$ speed advantage over the ego vehicle. Around a longitudinal position of $150m$ (Figure 4A - Human), the ego vehicle initiates a lane change to overtake the slower lead vehicle. After the lane change is initiated, the left-following vehicle starts braking to align with the speed of the ego vehicle.

The model simulation reproduces these elements of behaviour (Figure 4A - Model). When the ego closes in on the slower lead, and enough evidence accumulates (denoted by the cross), the benefit-based trigger initiates a lane change. Similar to the human driver, the ego does not attempt a strong acceleration upon moving into the left lane. Instead, it continues at roughly the same speed it had before leaving the right lane. This causes the model's left-following vehicle to perceive a high risk, prompting it to decelerate and align with the ego's slower speed.

However, there is a slight difference in the timing of the braking of the left-following vehicle. In both the human behaviour and the model simulation, the ego roughly initiates the lane change at the same time. However, the left-following vehicle only starts decelerating after the ego almost crosses the lane boundary (Figure 4A - longitudinal velocity). Meanwhile, for human behaviour, the left-following vehicle anticipates the ego's speed right after the lane change initiation. This difference arises primarily due to the model's simplified belief construction of other agents, which is based on extrapolating current positions, velocities, and accelerations rather than inferring intent early on.

4.2.2. Scenario B: Low-Speed-Difference. In scenario B (Figure 4B), the same three vehicles travel on a two-lane highway, but no lane change is initiated. The key interaction is between the ego and the lead vehicle, which stays in the same lane throughout. The ego vehicle follows closely, adjusting speed to any fluctuations in the leader. This results in the characteristic *accordion effect*. As the follower accelerates more quickly or the leader slows down, the gap between them shrinks. Conversely, the gap expands again when the follower reduces acceleration, or the leader regains speed.

In the human data, the ego decelerates while closing the gap to the lead vehicle before accelerating again to maintain a safe following distance. The model simulation reproduces these behaviour characteristics (Figure 4B - Longitudinal velocity). The ego vehicle decelerates until it reaches a speed below that of the lead vehicle, then accelerates slightly due to a lower bound replan, causing the gap to shrink. Once its perceived risk threshold rises, an upper-risk threshold replan restores a more comfortable distance (Figure 4B - Perceived risk). A notable difference is that the model holds a slightly larger preferred distance than the human driver, causing the gap to open and close one additional time (Figure 4B - Gap).

This lane-keeping behaviour is not explicitly programmed, and the planner does not impose any cost related to gap-keeping terms. Instead, these distance gaps occur from the interaction between risk perception and the other driver's probabilistic belief, which serves as a constraint in the plan optimization process.

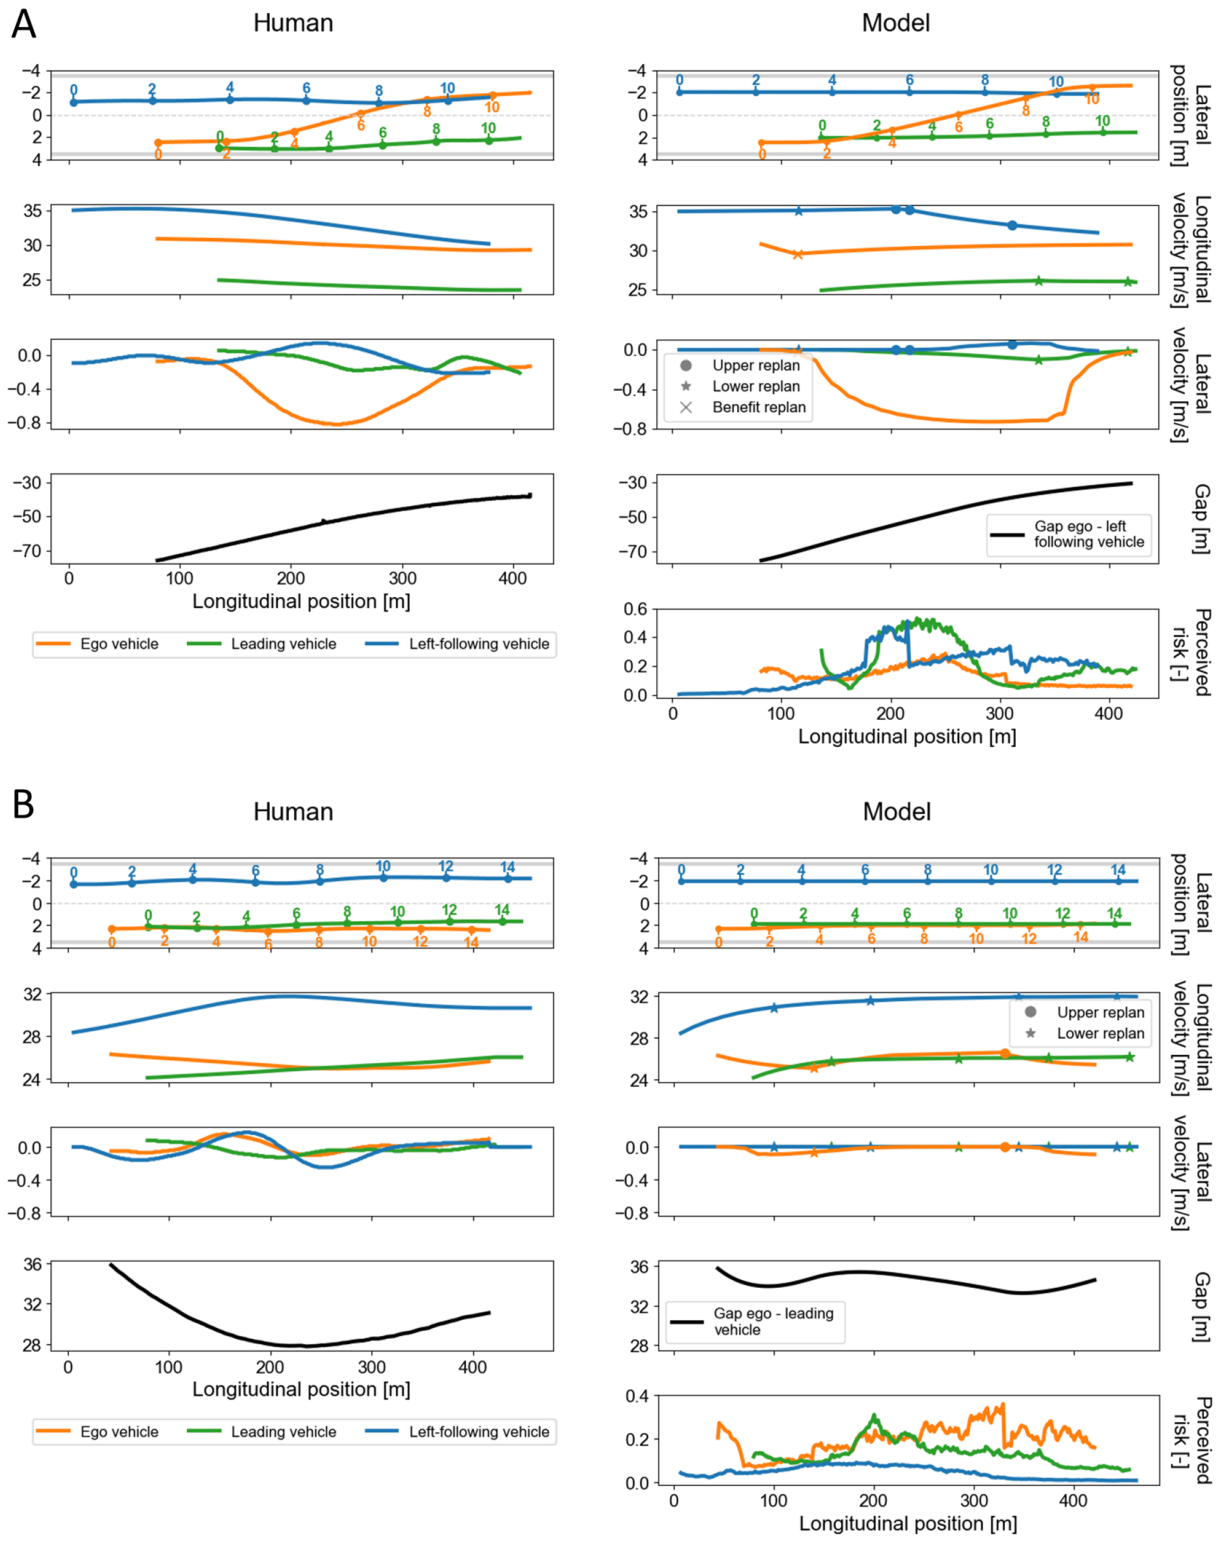

Figure 4: An overview of the scenario-level validation. **A)** Human and model behaviour in scenario A (high-speed-difference). The ego (orange), leading (green), and left-following (blue) vehicles correspond to vehicles 952, 949, and 955, respectively, in Figure 3A. The first row shows vehicle positions over time, with numbers indicating positions at 2.0 s intervals. The second and third rows present longitudinal and lateral velocities, where large dots indicate replanning due to exceeding the upper-risk threshold, stars for falling below the lower-risk threshold for longer than the saturation time, and crosses for a benefit-based replan. The fourth row depicts the gap between the ego and the left-following vehicle relative to the lead vehicle's position. The final row illustrates perceived risk over time. **B)** Human and model behaviour in scenario B (low-speed-difference). The ego (orange), leading (green), and left-following (blue) vehicles correspond to vehicles 53, 51, and 54, respectively, in Figure 3B.

4.3. Tactical Behaviour Validation

Following Siebinga *et al.* [38], I classified tactical behaviours into lane-keeping, lane-changing, off-road driving, and collisions, with the first two considered desirable and the latter two undesirable.

In scenario A (high-speed-difference), both the model and human ego vehicle exhibit similar behaviour across the 3000 simulated samples: the lane-change rate is 36.4% for the model and 36.0% for the human drivers, with corresponding lane-keeping rates of 63.7% and 64.0% (Table 1). This close alignment indicates that the model reproduces human decision-making when the speed differential is significant, providing a clear incentive to change lanes.

In scenario B (low-speed-difference), the incentive to change lanes is reduced, and as expected, both the model and human drivers predominantly opt to keep their lanes (Table 1). Here, the lane-keeping rates are 77.5% for the model and 86.0% for human drivers. However, the model initiates lane changes slightly more often (22.4% compared to 15.0% in the human data), suggesting that it may be more sensitive to lane-change opportunities in conditions where the speed differential is minimal. Still, the overall trend remains consistent; lane-keeping is the more prevalent decision over all samples in scenario B because of the small speed differential to the lead vehicle.

In addition to the results on desired behaviour, the results on undesired behaviour indicate a robust safety performance: no collisions occurred in any scenario, and off-road events were minimal, with only 0.1% occurrence in scenario A ($n = 4$) and 0.1% in scenario B ($n = 2$).

The strong overlap in both scenarios and the reasonable split between lane-changing and lane-keeping provide evidence that the proposed approach reproduces the discretionary characteristics of tactical manoeuvres while reflecting the heterogeneity in the human data. While the goal of this model is not to predict specific lane-change or lane-keeping samples, a one-to-one comparison (Table 2) is performed to prevent misleading conclusions. This ensures that a well-matched overall distribution does not mask discrepancies.

In scenario A (high-speed-difference), Table 2 shows that when human drivers chose to change lanes, the model did so in 79.8% of these samples, whereas for human lane-keeping decisions, the model's corresponding match was 87.8%. In scenario B (low-speed-difference), Table 2 shows that the model slightly overestimates lane-change actions. Specifically, among the samples where humans opted for a lane change, the model chose lane change 83.3% of the time, leaving a higher mismatch rate of 11.7% for lane-keeping decisions in these samples.

Figure 5 illustrates the ego's lane-changing and lane-keeping responses over the course of scenarios A and B, providing a qualitative perspective on tactical behaviour. It visualises how behaviour unfolds across time for both the human data and the model in one representative iteration of the 10 runs per scenario.

In scenario A, many human drivers initiate a lane change between $50 - 200m$ of travel. The model exhibits a similar split, with the ego vehicle often changing lanes at roughly the same positions. As shown by the lane change decision rate in Table 1, the proportion of lane changes is closely matched, indicating that the model captures both quantitative and qualitative aspects of tactical decision-making. By contrast, in scenario B (low-speed-difference), both humans and the model predominantly remain in their lane. Consequently, lane-keeping appears as the most common decision, driven by the reduced speed differential with the lead vehicle.

Beyond these overall trends, the model also shows some emergent behaviour. In three samples (two in scenario A and one in scenario B), the model performed a second lane change immediately after the first (Figure 5 -

Table 1: Tactical behaviour of the ego vehicle in scenarios A (high-speed-difference) and B (low-speed-difference), comparing lane-change, lane-keeping, off-road events, and collision rates between the model and human drivers. Each scenario contained 300 samples and was repeated 10 times, yielding 6000 simulations in total. The reported percentages represent the mean and standard deviation across these repeated simulations.

	Scenario A: high-speed-difference		**Scenario B: low-speed-difference**	
	Percentage of model behaviour	Percentage of humans behaviour	Percentage of model behaviour	Percentage of humans behaviour
Lane-change	36.2 ± 1.2	36.0	22.4 ± 2.7	15.0
Lane-keeping	63.7 ± 1.2	64.0	77.5 ± 2.8	85.0
Off-road	0.1 ± 0.2	0.0	0.1 ± 0.1	0.0
Collision	0.0 ± 0.0	0.0	0.0 ± 0.0	0.0

Table 2: One-to-one comparison of lane-change and lane-keeping decisions between human and the model ego vehicles in scenarios A and B. The table presents absolute counts of the model's decisions (rows) given the human driver's decision (columns), with column-wise percentages in parentheses. For scenario A, four off-road samples are excluded, and for scenario B, two off-road samples are excluded.

	Scenario A: high-speed-difference		**Scenario B: low-speed-difference**	
	Human lane-change	Human lane-keeping	Human lane-change	Human lane-keeping
Model lane-change	852 (79.8%)	236 (12.2%)	373 (83.3%)	298 (11.7%)
Model lane-keeping	216 (20.2%)	1692 (87.8%)	75 (16.7%)	2252 (88.3%)

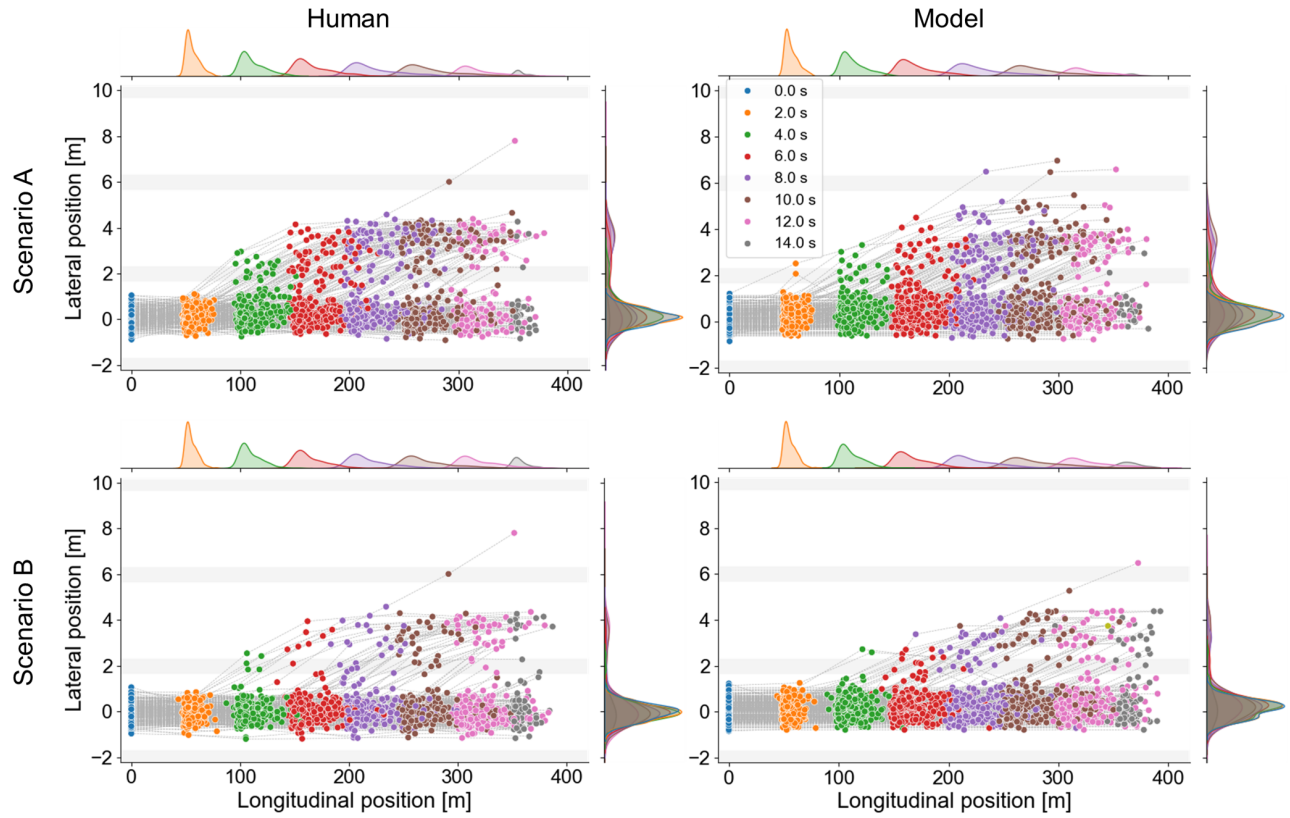

Figure 5: An overview of the variability in ego's responses for the human and model behaviour in scenarios A and B. This figure shows results for one representative iteration of those 10 runs per scenario to keep the presentation comprehensible. The top row represents scenario A, while the bottom row represents scenario B. The left column corresponds to human behaviour, and the right column represents the model behaviour. The longitudinal and lateral positions are normalized so that $0m$ represents the longitudinal position on the highway in the first frame and the centre of the right-most driving lane. Each dot in the figure corresponds to one time step in one of the 300 samples, and the distributions at the top and right sides of the figure show how the variability in response spreads out over time. The horizontal grey bars denote the span in which all lane markings are positioned. Operational variability can also be observed in human and model behaviour regarding longitudinal and lateral positions for lane-keeping and lane-changing. This figure uses adaptations from [22].

Model). Typically, vehicles accelerate after an initial lane change to create a safe gap; however, these particular drivers opted for an additional lane change to mitigate high risk due to a faster-approaching vehicle from behind. Because the current model does not explicitly enforce traffic rules and norms, this second lane change emerges from the balance between steering and acceleration plan weight, where the cost of a second lane change was lower than accelerating to reduce risk.

Moreover, three agents aborted their lane changes and returned to the rightmost lane, one for scenario A (at $2.0s$, around $50m$) and two for scenario B (at $4.0s$, around $125m$). This demonstrates the model's ability to handle "failed" lane change attempts. Although sufficient space was initially available, the increasing risk from an approaching following vehicle triggered a replan back to the rightmost lane.

4.4. Operational Behaviour Validation

The following validation aggregates data from scenarios A and B but only includes those samples in which both the model and human drivers exhibited the same tactical behaviour (i.e., correctly reproduced samples from Table 2). The figures in this subsection show results for one representative iteration of those 10 runs per scenario to keep the presentation comprehensible.

4.4.1. Distinct Execution Characteristics from Different Tactical Choices. Figure 6 compares the model's variability in operational behaviour for different tactical behaviours (lane-changing and lane-keeping) with human behaviour using inverse time-to-collision (TTC) vs time gap. In the human data, during a lane change (Figure 6 - Lane-change), the inverse TTC increases while the time gap decreases until the moment the vehicle crosses the centre lane marking, represented by the green circles. This increase in inverse TTC reflects a typical human tendency, also noted in empirical research, to accelerate and close the gap to the lead vehicle during a lane change to overtake a slower lead vehicle [30, 39]. Similarly, the model's lane-changing behaviour exhibits dynamics that closely resemble those observed in human behaviour. However, the model makes some lane changes at a lower inverse TTC (higher TTC) than humans.

When comparing the model's lane-keeping behaviour to that of human drivers, oscillations around an equilibrium point are observed for most samples (see Figure 6, lane-keeping). Empirical studies suggest that human lane-keeping can be characterized by cyclic behaviour around an approximate equilibrium spacing [38, 40], as seen in the inverse TTC vs time gap plot. These oscillations reflect the adjustments drivers make to maintain a stable following distance.

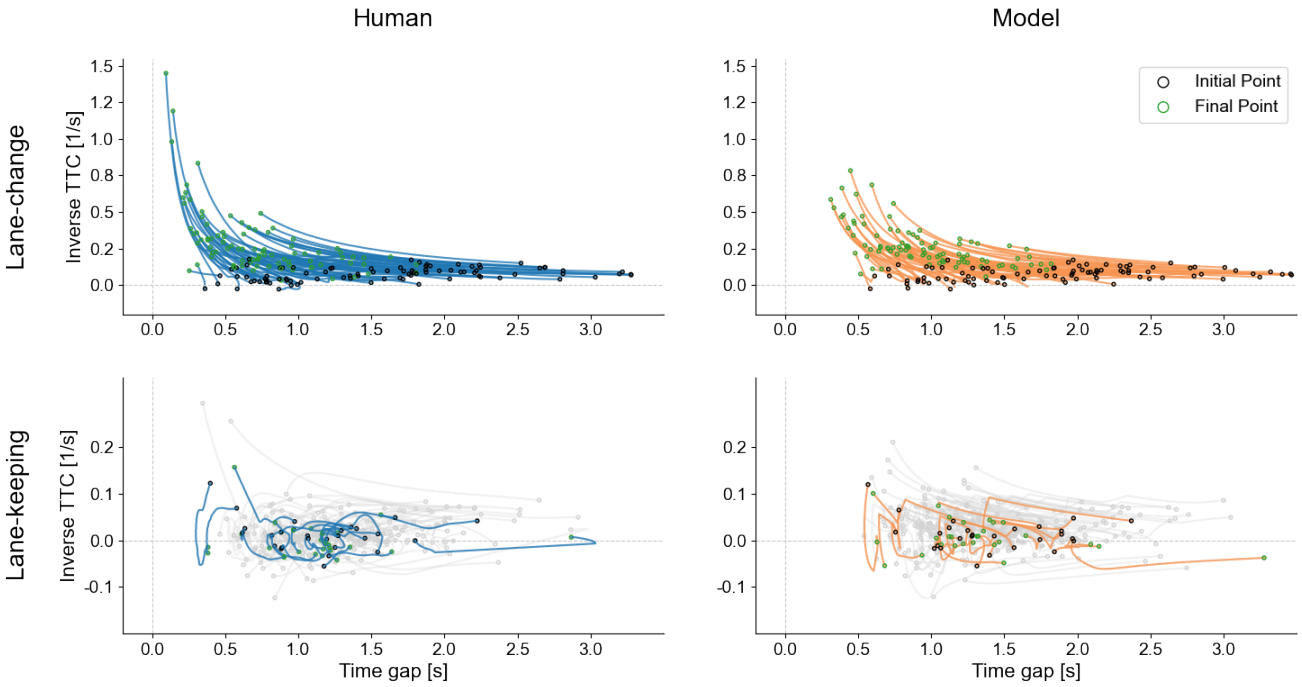

Figure 6: An overview of the operational characteristics for different tactical behaviours. Blue lines denote the human behaviour and orange lines the model behaviour. The first row shows lane-changing behaviour, while the second row depicts lane-keeping behaviour. Black circles represent the initial positions of the ego vehicle, and green circles mark the final positions. For lane-changing samples, the green circles indicate the moment of boundary crossing at the centre-line between lanes. To facilitate a clear comparison for lane-keeping, 25 random sample pairs are highlighted. This figure uses adaptations from [38].

Compared to human drivers, the model decelerates more abruptly when the time gap decreases upon reaching its upper-risk threshold (see Figure 6 - Lane-keeping). This leads to sharper "corners" in the model's lane-keeping behaviour. In contrast, observations from the human data indicate that drivers typically adjust their speed in smaller increments, leading to smoother transitions. After slowing down, the model also tends to maintain a larger gap for a longer period, whereas human drivers usually accelerate sooner once the gap has sufficiently widened. Consequently, human drivers exhibit higher-frequency oscillations around the equilibrium compared to the model. These differences can be partly explained by the model's discrete replan intervals. Following the initial replan triggered by risk, there is often a prolonged delay before the next replan occurs (e.g., upon crossing the preferred velocity or reaching the saturation time). Once triggered, this second replan ensures that the model no longer remains in an overly cautious state once the conflict is resolved.

However, overall, the model consistently reproduces distinct operational characteristics that emerge from different tactical decisions, aligning with the patterns observed in human lane-changing and lane-keeping.

4.4.2. Accepted Gaps at Lane Change Initiation. To further validate the operational lane change behaviour, I investigated the distributions of time gap at the moment of lane-change initiation for all surrounding vehicles from the ego vehicle's perspective (Figure 7A-B). Across all interactions, the model closely reproduces the variability in human time gap distributions. For the ego–lead vehicle time gap, the medians of the human and model overlap at approximately $1.1s$. The time gap between the ego and both left-lane vehicles also shows broad consistency, with the model reproducing the primary region where drivers accept gaps in the target lane. In the corresponding scatter plots (Figure 7B), the modelled time gaps correlate strongly with measured human values (R^2 between about 0.7 and 0.8), and the regression slopes are close to unity.

By contrast, when examining inverse TTC at lane change initiation (Figure 7C-D), the distributions still overlap but differ more in shape and central tendency. The scatter plots confirm that inverse TTC from the model aligns less consistently with individual human-measured values (Figure 7D), with lower R^2 scores than observed for the time gap. These differences are likely due to the model's assumption of constant desired velocity. In the human data, the speed of surrounding vehicles is influenced by other vehicles in the traffic stream, leading to occasional under- or overestimates of how quickly a gap is closing. Nevertheless, even with these discrepancies in inverse TTC, the model generally reproduces the range of values drivers accept at lane-change onset, suggesting that the overall gaps for lane-changing are similar to human behaviour despite the simplified velocity assumptions.

4.4.3. Speed Adjustments During Lane Changes. After accepting a gap, drivers proceed with the execution of the lane change manoeuvre. Empirical evidence shows that when overtaking a slower lead vehicle, drivers often decelerate gradually before the lane change to avoid colliding with the slower vehicle in front [30, 39]. After the lane change onset, drivers accelerate to the passing speed to adopt a speed according to the speed of the vehicles in

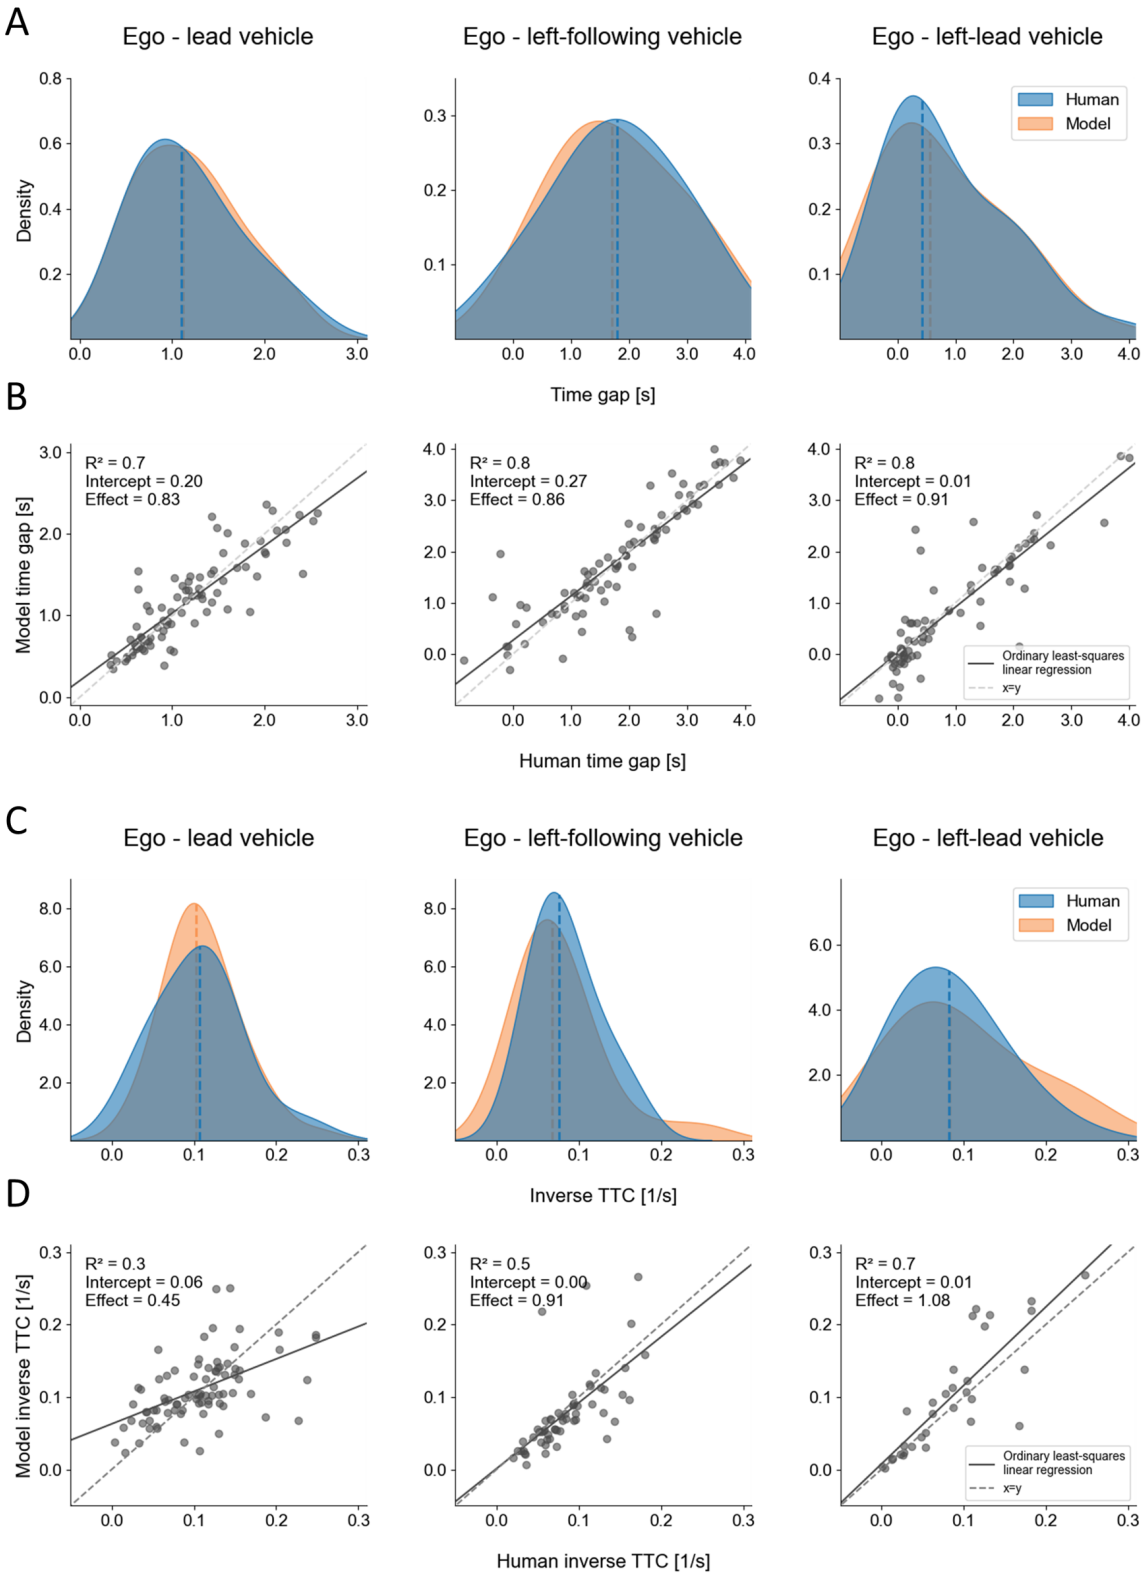

Figure 7: An overview of the accepted gaps at lane change initiation. The blue distributions denote the human behaviour and the orange distributions represent the model behaviour. **A)** Estimated distributions of time gap at the moment of lane change initiation to the lead, left-following, and left-leading vehicle. Dashed lines indicate the medians of each distribution. **B)** The relationships between the human and model behaviour for all data points corresponding to the time gap distributions. **C)** Estimated distributions of inverse TTC at the moment of lane change initiation to the lead, left-following, and left-leading vehicle. The inverse TTC axis is capped between -0.05 and 0.3 to exclude extremely high values (i.e., TTC close to zero), preventing dominance by the broad distribution, especially for the ego-left lead vehicle. **D)** The relationships between the human and model behaviour for all data points corresponding to the inverse TTC distributions. Dashed lines indicate the medians of each distribution.

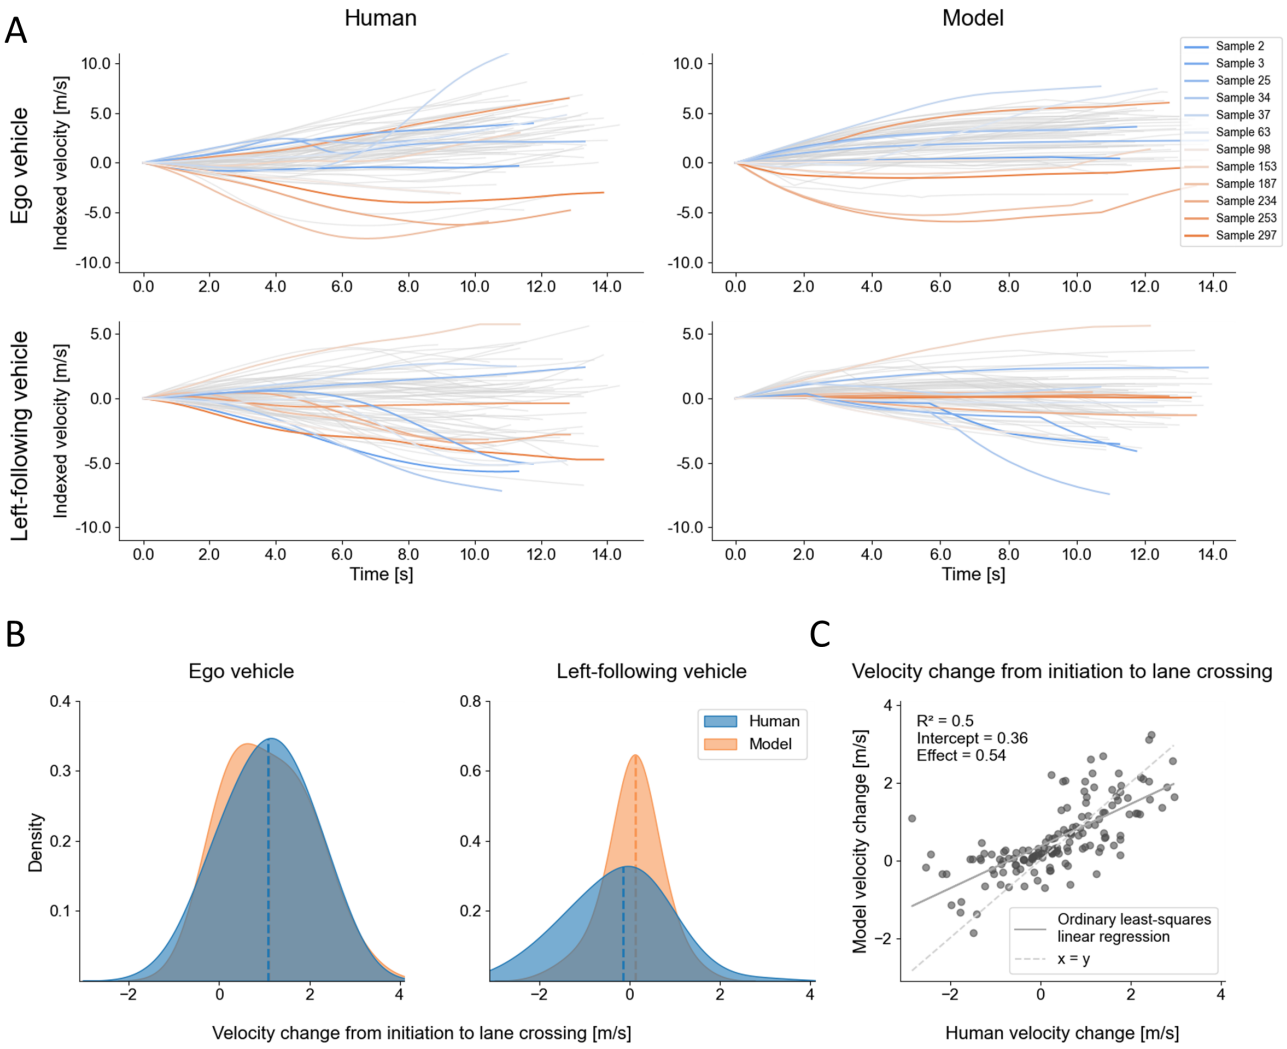

Figure 8: An overview of the control inputs performed by the human drivers and by the model for the ego vehicle performing the lane change and the left-following vehicle. **A)** All velocity traces where the ego performed a lane change for the human driver (first column) and the model (second column). All velocities are indexed based on their initial velocity. Twelve random pairs are highlighted to facilitate a clear comparison of individual samples. **B)** The change in velocity from lane change initiation to crossing the lane boundary for the ego vehicle and the left-following vehicle responding to the ego's lane change. Dashed lines indicate the medians of each distribution. **C)** The relationship between human and model behaviour for both the ego and left-following vehicle, corresponding to B).

the target lane [30, 39]. This results in considerable speed adjustments throughout the manoeuvre.

This pattern is also visible in the human data (Figure 8A - Ego vehicle), where drivers undertake substantial speed changes by gradually reducing speed before committing to the lane change and then accelerating once it is underway. For the most part, the model reproduces these observed patterns, exhibiting a "deceleration-then-acceleration" sequence that is qualitatively similar to human behaviour. Additionally, both the human and model data also show that a substantial number of samples involve drivers accelerating immediately and changing lanes without a prior deceleration phase, suggesting an alternative approach to executing the manoeuvre to prevent any speed loss.

Nonetheless, the model occasionally displays a more pronounced or prolonged deceleration than humans. Since both positive and negative deviations from a desired speed are penalized equally, the model sometimes chooses to keep slowing down until it becomes clear that this lower speed is no longer optimal. Only then does it switch to accelerating and complete the lane change. This happens partly because the model does not continuously optimize its plan. It first lowers its speed to reduce risk, then waits until it sees that staying slow is also risky or inefficient, prompting the lane change and a switch to a higher speed. In contrast, observations from naturalistic data indicate that human drivers typically blend these steps and move sooner into the faster lane.

Another discrepancy arises with the response of the left-following vehicle (Figure 8A - Left-following vehicle). Human drivers anticipate lane changes of other vehicles using subtle implicit cues, like slight lateral movements or speed changes, and explicit communication, such as indi-

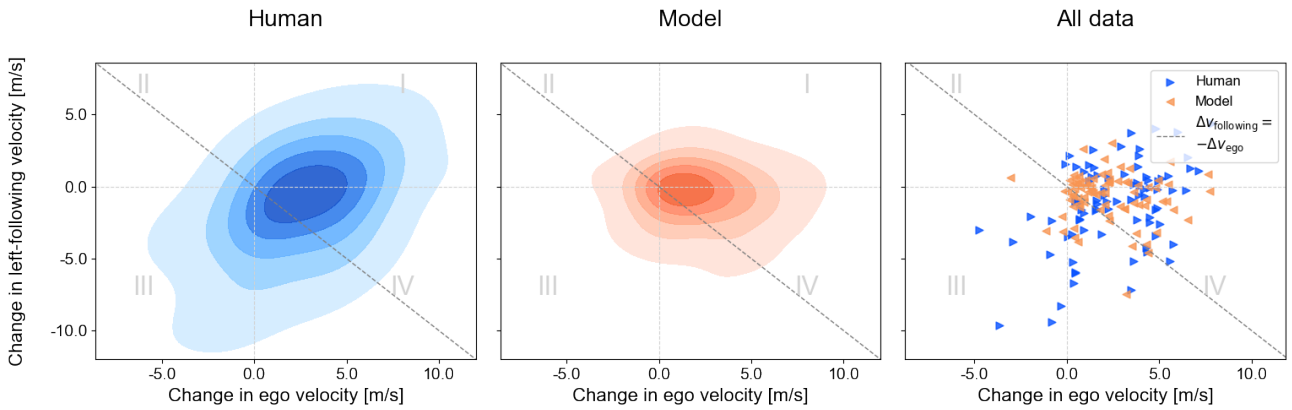

Figure 9: An overview of the velocity contributions to preserve or enlarge the gap of the ego (x-axis) and left-following (y-axis) vehicles after lane change initiation, with velocities indexed at zero at the initiation. The dashed diagonal line indicates contributions ($\Delta v_{follower} = -\Delta v_{ego}$), where both vehicles' adjustments equally impact the distance gap. The figure is divided into quadrants that indicate whether each vehicle contributes positively or negatively to maintaining a safe gap. Quadrant IV represents positive contributions to preserving or enlarging the gap (ego accelerates, left-following decelerates), while quadrant II shows both vehicles' negative contributions (reducing the space gap).

cator lights [41]. By contrast, the model's left-following vehicle only reacts once it has detected a more definitive start to the lateral manoeuvre. Its beliefs rest on extrapolating current positions, velocities, and accelerations rather than inferring intent early on. As a result, the model underestimates how forcefully or courteously other drivers respond to an upcoming lane change, contributing to smaller velocity deviations compared to human behaviour (Figure 8A - Left-following vehicle).

The estimated distributions of change in velocity from lane change initiation to lane centre crossing (Figure 8B) further support these observations. The ego vehicle's distribution aligns well between human and model data, indicating that the model reproduces the expected speed changes during lane changes. However, the left-following vehicle in the model exhibits a more constrained distribution, with smaller velocity adjustments compared to human drivers, who demonstrate a broader range of speed adaptations. This discrepancy likely stems from the model's simplified belief, as the speed adjustments of the left-following vehicle remain minimal until the moment the ego vehicle crosses into its lane. This difference is also reflected in the regression analysis (Figure 8C). The correlation between modelled and human relative velocity is moderate ($R^2 = 0.5$), suggesting that while the model captures the general trend, it underestimates the extent of speed adaptations.

Through their individual control inputs, the ego and left-following vehicle collectively contribute to maintaining a safe gap, which can be viewed as gap-keeping: maintaining or adjusting the distance to the newly leading or newly following vehicle. Figure 9 shows how the ego and the left-following vehicles adjust their speeds from the start to the end of the lane-change manoeuvre, showing each relative contribution to preserving or enlarging the accepted gap after the lane change initiation.

Based on observed human behaviour, the ego vehicle promptly accelerates to a passing speed for most samples (quadrant I and IV in Figure 9 - Human). The left-following vehicle decelerates most of the time to a lower speed (quadrant III and IV in Figure 9 - Human), further contributing to maintaining a safe gap. For the ego's behaviour in the model, the vehicle tends to increase its velocity most of the time, thereby contributing to preserving or enlarging the gap (quadrant I and IV in Figure 9 - Model). However, as noted earlier, the model underestimates how much the left-following driver anticipates and reacts to the lane change. As a result, the model's left-following vehicle contributes less to maintaining or enlarging the gap compared to the ego vehicle.

Despite these slight underestimations in the left follower's braking, the model's overall gap-keeping behaviour aligns reasonably well with human patterns. Both human drivers and the model exhibit the "accelerate-decelerate" pattern for the ego and left-following vehicle. Minor anticipatory discrepancies arise from the model's simplified belief and risk components.

5. Discussion

In this work, I have presented a computational model for discretionary lane change interactions, integrating decision-making and low-level control into a unified model based on the CEI framework [25]. The extension of this framework to a two-dimensional lane change model demonstrates its ability to reproduce both tactical (i.e., *which* manoeuvres are executed) and operational (i.e., *how* the manoeuvres are executed) aspects of driving.

The model reproduces realistic human-like tactical variability in response to a typical discretionary lane change scenario, where drivers either keep their lane or change lanes to overtake a slower lead vehicle. It successfully captures variability in operational behaviour, showing that different tactical choices lead to distinct, physically meaningful execution characteristics. Furthermore, it quantitatively and qualitatively reproduces lane-changing characteristics and reveals lane-keeping behaviour consistent with empirical observations. Notably, these effects emerged

despite the cost function not having an explicit time gap or distance gap term.

Based on communication-based belief and risk- and benefit-based replanning, the model can dynamically balance multiple driving objectives (e.g., maintaining safety margins or achieving a desired speed) while adapting its behaviour in response to interactions with multiple surrounding vehicles. These findings demonstrate that the proposed unified model offers a promising and novel approach for modelling reciprocal multi-agent interactions.

5.1. Relation to the Existing Literature

One of the most prevalent approaches for modelling lane change interactions is game theory (e.g., [10, 11]). Similar to game theory, the findings suggest that incorporating multiple drivers within a single model (as opposed to modelling individual drivers) is important, as collective behaviours emerge from the combined actions of individual drivers. In contrast with game theory, the results demonstrate that the proposed model is not limited to a discrete set of actions (e.g. accelerating change lane, decelerating change lane, or doing nothing [10]). Instead, it models variability in tactical and operational behaviours in response to interactions with surrounding traffic.

Several existing single-agent models for drivers' lane-changing motivations rely on instantaneous decision-making processes (e.g., [4, 5, 12, 13]), thus overlooking how humans dynamically process information over time. Although these models are useful for large-scale microscopic traffic analyses, they do not reflect how drivers interpret and accumulate evidence before changing lanes. In contrast, the proposed model integrates implicit communication and beliefs about other's intentions to accumulate evidence and make decisions more human-like. Nonetheless, this study adopts a simplified incentive, assuming that drivers primarily change lanes to gain a speed advantage based on relative speed and distance to surrounding traffic. A more sophisticated approach could also consider factors such as traffic density or travel times in adjacent lanes [19, 42].

Furthermore, most microscopic simulation models rely on the same gap acceptance theory (e.g., [5, 12, 13]), which states that drivers accept any gap larger than a critical threshold. However, empirical research reported that drivers sometimes reject gaps exceeding this threshold or accept gaps lower than the critical gap [43, 44], indicating inconsistencies in the theory. Unlike gap acceptance theory, the proposed approach is not assumed to have a binary decision (accept/reject) when a driver commits to a lane change. The proposed approach models gap acceptance through the risk of a probability of a collision, using probabilistic beliefs about other drivers' future actions derived from implicit communication. Hence, gap acceptance is based on the reciprocal interactions between drivers, and there is no fixed critical gap; drivers adapt to each situation based on contextual risk. This enables the model to represent failed lane change attempts, a frequent real-world occurrence that existing lane-changing decision models fail to consider. These failed lane change attempts are not explicitly modelled, but they occur due to combinations in changing communication and subsequent belief and risk perception. This study is among the first to report failed lane change attempts in discretionary lane change manoeuvres, and there is a need for further research to explore the underlying mechanisms leading to these changes in decision-making.

Finally, the presented model builds upon the risk-based replanning of the CEI framework, a mechanism introduced previously by Kolekar *et al.* [45]. In Kolekar's model, the perceived risk is derived by multiplying the cost of an event (e.g., overtaking, going off-road, or encountering obstacles) by its associated *Driver Risk Field*. I adopted a simplified definition of these risk fields, focusing on maintaining a safe distance within lane boundaries and preventing off-road travel. A key distinction shared by both approaches is that neither searches for an optimal solution continuously. Instead, each model seeks to achieve a specific objective (such as maintaining a desired speed) while keeping the risk below a personal threshold. Humans, for instance, do not optimise to drive perfectly in the centre of a lane but rather stay within acceptable lateral limits. The model reproduces this variability by using perceived risk to maintain position within these limits, similar to Kolekar's Driver Risk Field.

5.2. Implications

The proposed model and the broader CEI framework have several implications for the development and validation of AVs. Human-like reference models are essential for the virtual testing of AVs. Yet, current human driver reference models often lack sophisticated interactive elements and diverge significantly from real driver behaviours [15]. The proposed lane change model did not result in any crashes during the validation. This outcome highlights the limitation of current reference models, which, as noted by Olleja *et al.* [15], tend to overestimate crash occurrences in cases where human drivers would not have collided. However, it is important to note that the model is primarily validated on typical interactive behaviours with moderate to low severity rather than exceptional behaviours such as those leading to crashes and/or serious conflicts (see [46] for the definition of safety-related behaviour). Future research should assess whether the model can prevent near-crash events.

For a CEI-based model to serve as a benchmark for AV testing, it must also capture driver reaction times and the associated stimuli [15, 47]. While previous work with CEI has demonstrated the framework's ability to reproduce the time required for drivers to achieve safe outcomes [26], the current implementation is not explicitly validated on reaction time and stimulus perception as it is not possible to unambiguously define stimuli from the used naturalistic driving data. Engström *et al.* [47] addressed this issue by proposing the Non-Impaired Eyes ON reference model, which accounts for situational dependence in response timing. Integrating conceptual insights from this model could enhance belief construction within a CEI-based model, enabling both the dynamic definition of stimuli and more

realistic belief updates leading to evasive responses. The main advantage over existing benchmark models could be that the CEI-based model is developed specifically for interactive scenarios and thus might produce more realistic interactive behaviour.

Beyond describing individual driver behaviour characteristics in AV validation, the model also shows promise to enhance virtual validation environments for AVs or in microscopic traffic simulations. The results indicate that this model can reproduce realistic manoeuvres and failures. As in real life, traffic flow characteristics are strongly influenced by how lane changes are executed [18]. Therefore, incorporating delayed or failed lane change attempts can impact traffic flow efficiency and safety. Accounting for these execution processes could improve the realism of microscopic simulations. It is worth noting that the validation I conducted in the previous section involved at most four vehicles. In reality, more vehicles can interact during lane changes. Extending this model to large-scale simulations raises questions about how beliefs, risk, and communication cues scale when dozens or even hundreds of vehicles are involved. The question of how this model performs in large-scale simulations is left for future work. Despite these challenges, the promise of a unified framework for multi-agent traffic interactions holds considerable appeal for industrial applications and further academic exploration.

5.3. Limitations, Recommendation, and Future Work

While the current study demonstrates the potential of an integrated CEI-based approach, further refinements are needed. Both the model implementation and the validation method have important limitations.

First, this study used simplified components (plan, communication, belief, and risk) to demonstrate the CEI framework for a two-dimensional lane change model. The plan is defined by desired velocity, acceleration, heading, and steering input. The weighting factors for these parameters ($\lambda_1 - \lambda_5$) strongly influence both the incentive of gaining speed and subsequent lane change decision once perceived risk surpasses the upper threshold. At present, precisely tuning these weights is essential to achieve desirable tactical and operational behaviour. Although this approach has been validated for straight highway segments and shown to be robust to changes in these parameters, its performance in other scenarios remains unknown. For example, extending the model to account for road curvatures does not differentiate between the steering demands of following a curved road and those of executing a lane change.

A possible solution to account for this is adding additional penalties in the cost function for deviating from the current lane, which could help distinguish between necessary steering adjustments for road curvature and intentional lane-changing behaviour. Alternatively, because risk is enforced through an optimization constraint in the model's planning process, one could expand perceived risk to include event costs, as suggested by Kolekar *et al.* [45]. Doing so would enable a CEI-based model to handle discretionary lane changes and potentially other manoeuvres (e.g., road curvatures or obstacle avoidance) without requiring separate plans for distinct driving objectives. In this setup, the benefit-based trigger could be incorporated by assigning a minimal cost to the perceived risk function for the current lane rather than treating it as a separate plan parameter.

Using a more sophisticated definition of perceived risk (e.g., [45]), factors like curvature, lane endings, and on-/off-ramps would be accounted for, supporting discretionary and mandatory lane changes. A unified CEI-based framework that covers these manoeuvres would be an exciting step, revealing how reciprocal interactions vary by lane change type. Since this study already shows that risk-based replanning works for discretionary lane changes, it can also be adapted to other lane change motives, each involving different levels of risk-taking. Current literature acknowledges that mandatory and discretionary lane changes often involve distinct risk thresholds [48, 49]. Accounting for these differences would allow a single framework to handle various lane change types within a unified interactive model.

Another limitation of the current model implementation is the simplified beliefs, where driver state estimates are treated as independent Gaussian distributions based on extrapolating observed position, velocity, and heading. The results show that this approach restricts the model's ability to anticipate courtesy manoeuvres, such as yielding space for other vehicles to change lanes. Moreover, while drivers are modelled as a joint interactive system, beliefs are still treated as isolated per agent. In practice, seeing a fast-approaching vehicle behind another vehicle in the next lane might prompt drivers to anticipate a lane change and adjust their speed. In contrast, the model only accounts for the lane change once the manoeuvre has already begun or even later. The communication used to construct these beliefs relies solely on implicit cues, whereas human drivers also depend on explicit communication, such as turn indicators, to anticipate and adapt to a lane-changing vehicle [41].

Incorporating more realistic, multimodal belief constructions, such as trajectory prediction models like Trajectron++ [50], or extending the model with explicit communication could enable the model to account for more reciprocal behaviours (e.g., courtesy lane changes or prematurely braking). Another possibility could be to represent lateral behaviours using motion primitives, reflecting a more human-like way of reasoning and anticipating the behaviour of others. Rather than expecting a lane change by extrapolating position, velocity, and acceleration, a driver could infer lateral intentions through predefined motion primitives from the moment a behavioural shift is detected.

Finally, the highD dataset constrained the validation process in three ways. Because the highD dataset reconstructs velocity and acceleration from positions, insights into control inputs are limited. Therefore, this study was limited to comparing the relative velocities during lane change execution to prove the plausibility of the model's operational behaviour. In addition, the highD highway segments are short ($420m$), which restricts observation of

long-term behaviour and includes lane changes that start before or end after the dataset's boundaries. Moreover, each driver contributes only one trajectory in all extracted samples used for this study, with considerable variability in initial positions and speeds between samples, making it challenging to validate operational behaviour fully. Relating the initial kinematics of vehicles to specific lane change decisions (e.g., choosing to merge in front of or behind another vehicle) becomes difficult without direct measurements or repeated observations of the same driver's actions.

Nevertheless, these constraints did not undermine the main conclusions of this work. Even though highD's short segments and reconstructed velocity data limited direct comparisons of control inputs and long-term gap-keeping behaviour, they still captured the lane change phases necessary to evaluate whether the model can reproduce realistic characteristics. The model's demonstrated ability to produce plausible tactical decisions and operational characteristics that are both acknowledged in empirical research, as seen in the data, remained valid within highD's scope. For future studies, however, collecting a highD-like dataset that covers longer stretches of road (e.g., multiple kilometres) would improve validation. Such data could include full lane change trajectories, from the first indication of intent through completion, allowing insights into long-term gap-keeping behaviour and repeated observations of the same drivers. A possible intermediate step could be test-track experiments, where control inputs and repeated measurements of driver behaviour are recorded. This would allow the validation of the variability between drivers and the variability within a single human's behaviour.

6. Conclusion

In conclusion, this research tries to overcome the gap in discretionary lane change models by simulating reciprocal interactions and integrating decision-making and low-level control actions within a unified model. The proposed model for discretionary lane change interactions is based on the Communication-Enabled Interaction framework [25]. The findings demonstrate that the model successfully reproduces quantitatively and qualitatively individual behaviours and collective contributions while also demonstrating how varied tactical choices yield distinct, human-like operational outcomes. Subsequent research can build on my findings to refine human behaviour models further and strengthen the validation of automated vehicles in dynamic, reciprocal traffic contexts, thereby supporting safer, more efficient mobility systems.

References

[1] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, "A Survey of Autonomous Driving: *Common Practices and Emerging Technologies*," *IEEE Access*, pp. 58443–58469, 2020.

[2] F. Duarte and C. Ratti, "The Impact of Autonomous Vehicles on Cities: A Review," *Journal of Urban Technology*, vol. 25, pp. 3–18, Oct. 2018.

[3] A. Schieben, M. Wilbrink, C. Kettwich, R. Madigan, T. Louw, and N. Merat, "Designing the interaction of automated vehicles with other traffic participants: design considerations based on human needs and expectations," *Cognition, Technology & Work*, vol. 21, pp. 69–85, Feb. 2019.

[4] P. Gipps, "A model for the structure of lane-changing decisions," *Transportation Research Part B: Methodological*, vol. 20, pp. 403–414, Oct. 1986.

[5] T. Toledo, C. F. Choudhury, and M. E. Ben-Akiva, "Lane-Changing Model with Explicit Target Lane Choice," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1934, pp. 157–165, Jan. 2005.

[6] C. Wang, F. Guo, R. Yu, L. Wang, and Y. Zhang, "The Application of Driver Models in the Safety Assessment of Autonomous Vehicles: Perspectives, Insights, Prospects," *IEEE Transactions on Intelligent Vehicles*, vol. 9, pp. 2364–2381, Jan. 2024.

[7] N. Li, D. Oyler, M. Zhang, Y. Yildiz, I. Kolmanovsky, and A. Girard, "Game-Theoretic Modeling of Driver and Vehicle Interactions for Verification and Validation of Autonomous Vehicle Control Systems," Aug. 2016.

[8] R. Donà, K. Mattas, and B. Ciuffo, "Towards Bi-Dimensional driver models for automated driving system safety requirements: Validation of a kinematic model for evasive lane-change maneuvers," *IET Intelligent Transport Systems*, vol. 17, pp. 1784–1798, Sept. 2023.

[9] C. Wang, F. Guo, S. Zhao, Z. Zhu, and Y. Zhang, "Safety assessment for autonomous vehicles: A reference driver model for highway merging scenarios," *Accident Analysis & Prevention*, vol. 206, p. 107710, Oct. 2024.

[10] Y. Ali, Z. Zheng, M. M. Haque, M. Yildirimoglu, and S. Washington, "CLACD: A complete LAne-Changing decision modeling framework for the connected and traditional environments," *Transportation Research Part C: Emerging Technologies*, vol. 128, p. 103162, July 2021.

[11] A. Ji and D. Levinson, "Estimating the Social Gap With a Game Theory Model of Lane Changing," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, pp. 6320–6329, Oct. 2021.

[12] W. J. Schakel, V. L. Knoop, and B. Van Arem, "Integrated Lane Change Model with Relaxation and Synchronization," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2316, pp. 47–57, Jan. 2012.

[13] A. Kesting, M. Treiber, and D. Helbing, "General Lane-Changing Model MOBIL for Car-Following Models," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1999, pp. 86–94, Jan. 2007.

[14] K. Mattas, G. Albano, R. Donà, M. C. Galassi, R. Suarez-Bertoa, S. Vass, and B. Ciuffo, "Driver models for the definition of safety requirements of automated vehicles in international regulations. Application to motorway driving conditions," *Accident Analysis & Prevention*, vol. 174, p. 106743, Sept. 2022.

[15] P. Olleja, G. Markkula, and J. Bärgman, "Validation of human benchmark models for Automated Driving System approval: How competent and careful are they really?," June 2024.

[16] Y. Ali, F. Hussain, M. C. Bliemer, Z. Zheng, and M. M. Haque, "Predicting and explaining lane-changing behaviour using machine learning: A comparative study," *Transportation Research Part C: Emerging Technologies*, vol. 145, p. 103931, Dec. 2022.

[17] N. Montali, J. Lambert, P. Mougin, A. Kuefler, N. Rhinehart, M. Li, C. Gulino, T. Emrich, Z. Yang, S. Whiteson, B. White, and D. Anguelov, "The Waymo Open Sim Agents Challenge," Dec. 2023.

[18] M. Rahman, M. Chowdhury, Y. Xie, and Y. He, "Review of Microscopic Lane-Changing Models and Future Research Opportunities," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, pp. 1942–1956, Dec. 2013.

[19] P. Huang, L. Zhang, H. Chen, H. Ding, and J. Cao, "Event-triggered optimisation of overtaking decision-making strategy for autonomous driving on highway," *IET Intelligent Transport Systems*, vol. 16, pp. 1794–1808, Dec. 2022.

[20] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical Review E*, vol. 62, pp. 1805–1824, Aug. 2000.

[21] R. Keane and H. O. Gao, "A formulation of the relaxation phenomenon for lane changing dynamics in an arbitrary car following model," *Transportation Research Part C: Emerging Technologies*, vol. 125, p. 103081, Apr. 2021.

[22] O. Siebinga, A. Zgonnikov, and D. Abbink, "Uncovering Variability in Human Driving Behavior Through Automatic Extraction of Similar Traffic Scenes from Large Naturalistic Datasets," in *2023*

IEEE International Conference on Systems, Man, and Cybernetics (SMC), (Honolulu, Oahu, HI, USA), pp. 4790–4796, IEEE, Oct. 2023.

[23] J. A. Michon, "A Critical View of Driver Behavior Models: What Do We Know, What Should We Do?," in *Human Behavior and Traffic Safety* (L. Evans and R. C. Schwing, eds.), pp. 485–524, Boston, MA: Springer US, 1985.

[24] A. Ji and D. Levinson, "A review of game theory models of lane changing," *Transportmetrica A: Transport Science*, vol. 16, pp. 1628–1647, Jan. 2020.

[25] O. Siebinga, A. Zgonnikov, and D. A. Abbink, "Modelling communication-enabled traffic interactions," *Royal Society Open Science*, vol. 10, p. 230537, May 2023.

[26] O. Siebinga, A. Zgonnikov, and D. A. Abbink, "A model of dyadic merging interactions explains human drivers' behavior from control inputs to decisions," *PNAS Nexus*, vol. 3, p. pgae420, Oct. 2024.

[27] R. Rajamani, *Vehicle Dynamics and Control*. Mechanical Engineering Series, Boston, MA: Springer US, 2012.

[28] B. Jörges and L. R. Harris, "Object speed perception during lateral visual self-motion," *Attention, Perception, & Psychophysics*, vol. 84, pp. 25–46, Jan. 2022.

[29] A. J. Fath, M. Lind, and G. P. Bingham, "Perception of time to contact of slow- and fast-moving objects using monocular and binocular motion information," *Attention, Perception, & Psychophysics*, vol. 80, pp. 1584–1590, Aug. 2018.

[30] S. Moridpour, G. Rose, and M. Sarvi, "Effect of Surrounding Traffic Characteristics on Lane Changing Behavior," *Journal of Transportation Engineering*, vol. 136, pp. 973–985, Nov. 2010.

[31] G. Markkula, Y.-S. Lin, A. R. Srinivasan, J. Billington, M. Leonetti, A. H. Kalantari, Y. Yang, Y. M. Lee, R. Madigan, and N. Merat, "Explaining human interactions on the road by large-scale integration of computational psychological theory," *PNAS Nexus*, vol. 2, p. pgad163, June 2023.

[32] A. Zgonnikov, D. Abbink, and G. Markkula, "Should I Stay or Should I Go? Cognitive Modeling of Left-Turn Gap Acceptance Decisions in Human Drivers," *Human Factors: The Journal of the Human Factors and Ergonomics Society*, vol. 66, pp. 1399–1413, May 2024.

[33] S. H. Mohammad, H. Farah, and A. Zgonnikov, "Modeling Gap Acceptance in Overtaking: A Cognitive Process Approach," in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, (Bilbao, Spain), pp. 5925–5931, IEEE, Sept. 2023.

[34] A. Zgonnikov, N. Beckers, A. George, D. Abbink, and C. Jonker, "Nudging human drivers via implicit communication by automated vehicles: Empirical evidence and computational cognitive modeling," *International Journal of Human-Computer Studies*, vol. 185, p. 103224, May 2024.

[35] L. L. Hoberock, "A Survey of Longitudinal Acceleration Comfort Studies in Ground Transportation Vehicles," *Journal of Dynamic Systems, Measurement, and Control*, vol. 99, pp. 76–84, June 1977.

[36] R. Krajewski, J. Bock, L. Kloeker, and L. Eckstein, "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)*, (Maui, HI), pp. 2118–2125, IEEE, Nov. 2018.

[37] C. Hill, L. Elefteriadou, and A. Kondyli, "Exploratory Analysis of Lane Changing on Freeways Based on Driver Behavior," *Journal of Transportation Engineering*, vol. 141, p. 04014090, Apr. 2015.

[38] O. Siebinga, A. Zgonnikov, and D. Abbink, "A Human Factors Approach to Validating Driver Models for Interaction-aware Automated Vehicles," *ACM Transactions on Human-Robot Interaction*, vol. 11, pp. 1–21, Dec. 2022.

[39] D. D. Salvucci and A. Liu, "The time course of a lane change: Driver control and eye-movement behavior," *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 5, pp. 123–132, June 2002.

[40] M. Mulder, M. M. Van Paassen, M. Mulder, J. J. A. Pauwelussen, and D. A. Abbink, "Haptic car-following support with deceleration control," in *2009 IEEE International Conference on Systems, Man and Cybernetics*, (San Antonio, TX, USA), pp. 1686–1691, IEEE, Oct. 2009.

[41] T. Stoll, J. Imbsweiler, B. Deml, and M. Baumann, "Three Years CoInCar: What Cooperatively Interacting Cars Might Learn from Human Drivers," *IFAC-PapersOnLine*, vol. 52, no. 8, pp. 105–110, 2019.

[42] Y. Ali, Z. Zheng, M. Mazharul Haque, M. Yildirimoglu, and S. Washington, "Understanding the discretionary lane-changing behaviour in the connected environment," *Accident Analysis & Prevention*, vol. 137, p. 105463, Mar. 2020.

[43] F. Marczak, W. Daamen, and C. Buisson, "Merging behaviour: Empirical comparison between two sites and new theory development," *Transportation Research Part C: Emerging Technologies*, vol. 36, pp. 530–546, Nov. 2013.

[44] M. Yang, X. Wang, and M. Quddus, "Examining lane change gap acceptance, duration and impact using naturalistic driving data," *Transportation Research Part C: Emerging Technologies*, vol. 104, pp. 317–331, July 2019.

[45] S. Kolekar, J. De Winter, and D. Abbink, "Human-like driving behaviour emerges from a risk-based driver model," *Nature Communications*, vol. 11, p. 4850, Sept. 2020.

[46] A. Svensson and C. Hyden, "Estimating the severity of safety related behaviour," *Accident Analysis & Prevention*, vol. 38, pp. 379–385, Mar. 2006.

[47] J. Engström, S.-Y. Liu, A. Dinparastdjadid, and C. Simoiu, "Modeling road user response timing in naturalistic traffic conflicts: A surprise-based framework," *Accident Analysis & Prevention*, vol. 198, p. 107460, Apr. 2024.

[48] M. Vechione, E. Balal, and R. L. Cheu, "Comparisons of mandatory and discretionary lane changing behavior on freeways," *International Journal of Transportation Science and Technology*, vol. 7, pp. 124–136, June 2018.

[49] T. Pan, W. H. Lam, A. Sumalee, and R. Zhong, "Modeling the impacts of mandatory and discretionary lane-changing maneuvers," *Transportation Research Part C: Emerging Technologies*, vol. 68, pp. 403–424, July 2016.

[50] T. Salzmann, B. Ivanovic, P. Chakravarty, and M. Pavone, "Trajectron++: Dynamically-Feasible Trajectory Forecasting With Heterogeneous Data," Jan. 2021.

A

Model and Simulation Parameters

Table A.1: Model Parameters

Description	Parameter	Value
Simulation Parameters		
Simulation time step	dt	$0.05s$
Time horizon	T	$8s$
Belief frequency	f_b	$4Hz$
Communication Parameters		
Longitudinal velocity noise scaling	β_{long}	0.5
Lateral velocity noise scaling	β_{lat}	0.05
Perception update rate	α_c	$0.1s$
Belief Parameters		
Longitudinal comfortable acceleration	$a_{c,\text{long}}$	$1m/s^2$
Lateral comfortable acceleration	$a_{c,\text{lat}}$	$0.1m/s^2$
Belief scaling factor	ϕ	3
Memory length	T_m	$2s$
Heading convergence factor	η_{conv}	0.8
Risk Parameters		
Risk bounds	$\begin{bmatrix}\rho_l \\ \rho_u\end{bmatrix}$	$\begin{bmatrix}U(0.2,0.3) \\ U(0.4,0.5)\end{bmatrix}$
Risk incentives	$\begin{bmatrix}\lambda_{u,\text{dx}} \\ \lambda_{u,\text{dv}} \\ \lambda_{l,\text{dx}} \\ \lambda_{l,\text{dv}}\end{bmatrix}$	$\begin{bmatrix}0.002 \\ -0.02 \\ 0.002 \\ -0.02\end{bmatrix}$
Risk steepness	η	5
Middle risk heading std	$\sigma_{heading}$	$0.01rad$
Plan Parameters		
Weight longitudinal velocity	λ_1	$4e{-}1$
Weight lateral velocity	λ_2	$1e1$
Weight for heading deviation	λ_3	$1e3$
Weight for steering input	λ_4	$1e3$
Weight acceleration input	λ_5	5
Saturation time	τ	$3s$
Evidence Accumulation Parameters		
Drift rate	α_b	$U(0.02,0.04)$
Critical value for speed difference	θ_{crit}	$U(3.0,4.0)m/s$
Boundary	B	1

B

Preliminary Controlled Simulations

This appendix presents preliminary checks to ensure the model behaves consistently in various traffic scenarios and generates plausible tactical and operational behaviour.

In the preliminary simulations, I systematically vary the number of considered agents and the initial positions and speed of these vehicles to observe how all agents interact and adapt their behaviours based on surrounding vehicles. All simulations follow the parameter settings in Appendix A but use a constant critical value of 4.0 for the benefit-based trigger. This ensures the trigger is primarily activated at higher velocity differences from the lead vehicle, allowing a focused evaluation of the risk-based trigger at lower relative speeds. Additionally, each agent's desired velocity is set to its initial velocity. Each simulation is performed on a two-lane highway of $800m$ length and $3.5m$ lane width. All simulations are performed with CEI agents.

B.1. Two Agents: Lead Vehicle Interactions

Scenario description:

This scenario demonstrates how the initial gap and speed difference between an ego vehicle and a lead vehicle influence the ego's lane-change decisions and timing. In these scenarios (Figure B.1), I consider two vehicles on a two-lane road, where an orange ego vehicle approaches a slower green lead vehicle. In each sub-scenario (Figure B.1a-f), I vary two parameters: the initial gap between the ego and the lead ($80m$, $50m$, or $30m$) and the lead vehicle's speed at ($25m/s$, $27m/s$, or $29m/s$). The ego vehicle always starts at $30m/s$.

In scenarios a) and b), the gap is $80m$; in scenarios c) and d), the gap is $50m$; and in scenarios e) and f), the gap is $30m$. The only difference for each pair is that the lead vehicle's initial speed is $25m/s$ in the first case and 27 or $29m/s$ in the second.

In each sub-scenario, both vehicles are initially in the same lane. The ego vehicle's faster speed eventually causes it to close in on the leading vehicle, prompting a decision to either stay behind and slow down or change lanes to maintain the desired speed.

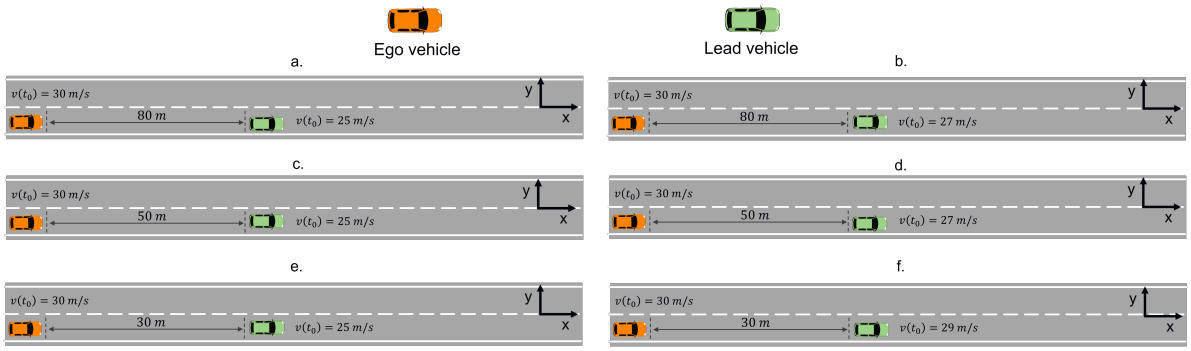

Figure B.1: Two-vehicle scenarios involving an orange ego vehicle and a slower green lead vehicle. The scenarios vary the initial gap between the vehicles ($80m$, $50m$, or $30m$) and the lead vehicle's speed ($25m/s$, $27m/s$, or $29m/s$).

Expected behaviour:

It is expected that as the initial gap between the ego and the lead vehicle decreases, the urgency of deciding whether to brake or overtake will increase. When the gap is large ($80m$), the ego vehicle should have sufficient time to react, allowing it to plan a gradual response, such as slowing down or making an early lane change to maintain the desired speed. At a medium gap ($50m$), the ego will likely encounter the lead vehicle sooner, requiring a faster decision to either brake or overtake before excessive speed loss occurs. At the smallest gap ($30m$), any speed difference should immediately prompt action of the ego vehicle, such as braking or a lane change, to prevent tailgating.

Similarly, changing the lead vehicle's speed from $25m/s$ to 27 or $29m/s$ reduces the speed difference with the ego, which is constant at $30m/s$. A large speed difference (e.g., $30m/s$ ego vs $25m/s$ lead) means the gap closes quickly, prompting the ego to brake earlier or plan an overtake. When the lead's speed is closer to the ego's speed (e.g., $29m/s$ vs $30m/s$), the gap closes more slowly or even stabilizes, so the incentive to overtake the lead vehicle is less. In practice, it is expected that a shorter initial gap and a slower lead together lead to a more decisive and earlier replan. In contrast, a more significant gap or a faster lead is expected to require only moderate adjustments from the ego vehicle.

Observed behaviour:

Across all six sub-scenarios (Figure B.2-B.7), the observed outcomes match the expected patterns of ego vehicle behaviour. Larger initial gaps and higher lead speeds induce slower closure rates, so the ego has more time to evaluate potential manoeuvres, often changing lanes later. Smaller gaps and greater speed differences, by contrast, trigger earlier decisions, with the ego either switching lanes almost immediately to avoid braking (as in sub-scenario c Figure B.4 and e Figure B.6) or moderating its speed quickly once the risk threshold is reached (as in sub-scenario b Figure B.3 and d Figure B.5). In sub-scenario f) (Figure B.7), the ego remains in the current lane and drops its speed below that of the lead vehicle to keep a safe following distance, then slightly accelerates to maintain the desired gap, demonstrating typical car-following behaviour.

This consistent interplay between the gap and speed difference demonstrates that the model behaves as expected. It balances the benefit of maintaining speed against the rising risk of tailgating and responds earlier when the threat is greater or the closure rate is high. Consequently, these simulations demonstrate that the ego vehicle's decision-making mechanisms function as designed.

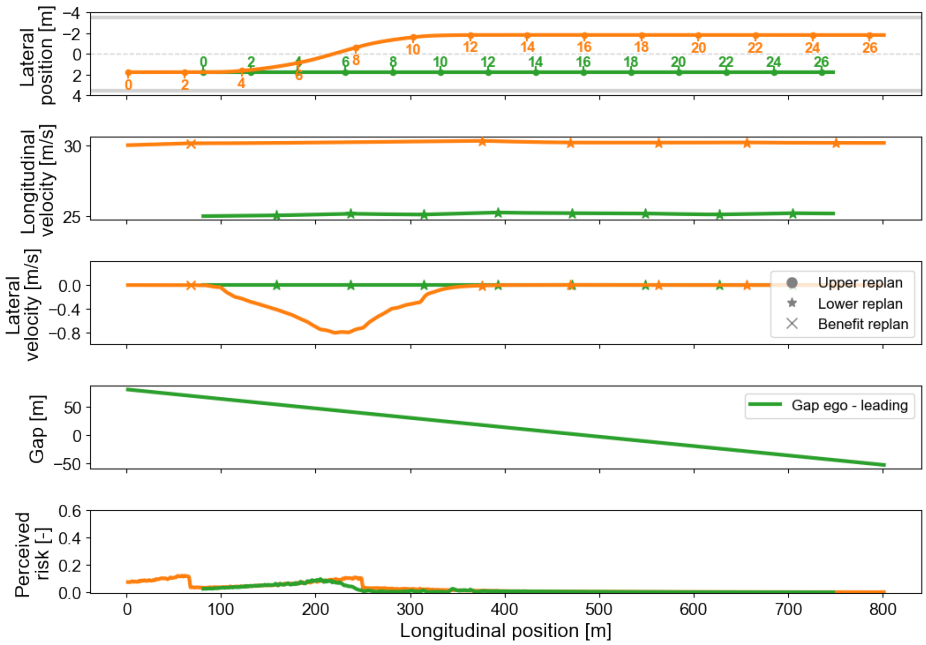

Figure B.2: Sub-scenario a). Initial conditions: the ego vehicle starts at $v(t_0) = 30$ m/s, with a gap of $80m$ to the lead vehicle, which has an initial speed of $v(t_0) = 25$ m/s. This large gap and speed difference creates a scenario where the ego vehicle has sufficient time to evaluate and execute its manoeuvre as it closes in on the lead vehicle. After $80m$ of travel, the benefit-based trigger initiates an early lane change.

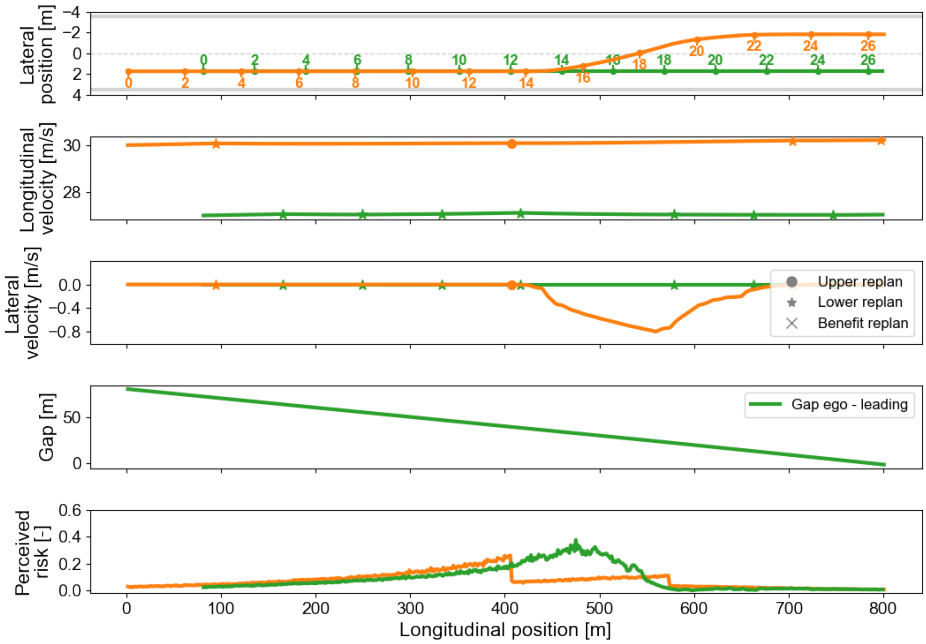

Figure B.3: Sub-scenario b). Initial conditions: the ego vehicle starts at $v(t_0) = 30$ m/s, with a gap of 80 m to the lead vehicle, which has an initial speed of $v(t_0) = 27$ m/s. The smaller speed difference relative to sub-scenario a) allows the gap to close at a slower rate, reducing the urgency of the ego vehicle's response. The speed difference is not high enough to trigger a benefit-based lane change; however, at a gap of $40m$, the upper-risk threshold is exceeded, and a lane change is initiated.

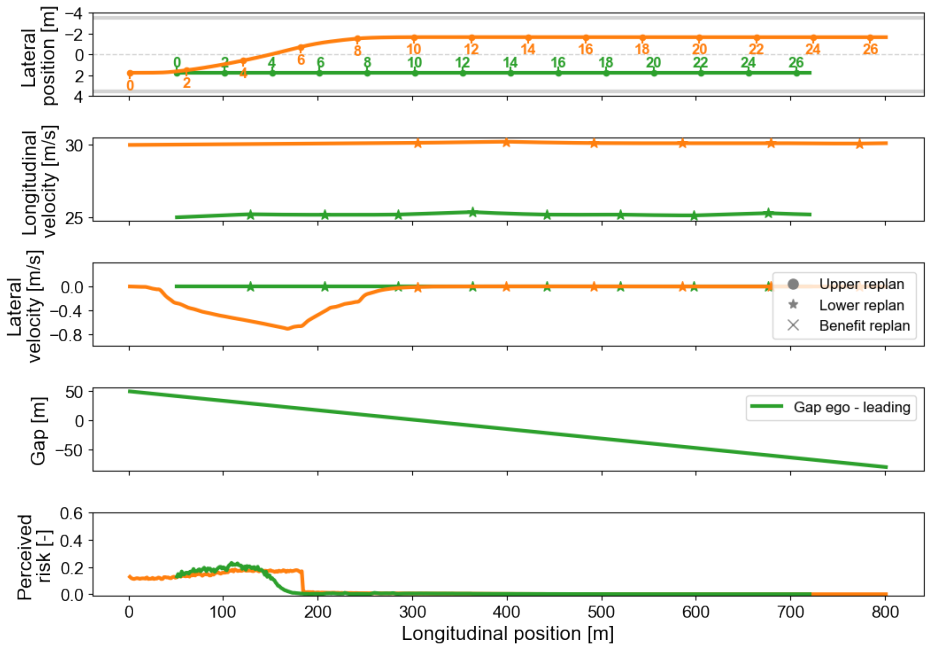

Figure B.4: Sub-scenario c). Initial conditions: the ego vehicle starts at $v(t_0) = 30$ m/s, with a gap of 50 m to the lead vehicle, which has an initial speed of $v(t_0) = 25$ m/s. With the smaller gap, the ego vehicle encounters the lead vehicle sooner, necessitating an earlier decision to either brake or overtake. The ego vehicle immediately plans a lane change at initialisation to prevent speed loss and reduce the perceived risk. During the manoeuvre, the risk remains high for the ego vehicle due to the lane proximity risk and insufficient heading adjustment, but it is resolved once the lane is crossed.

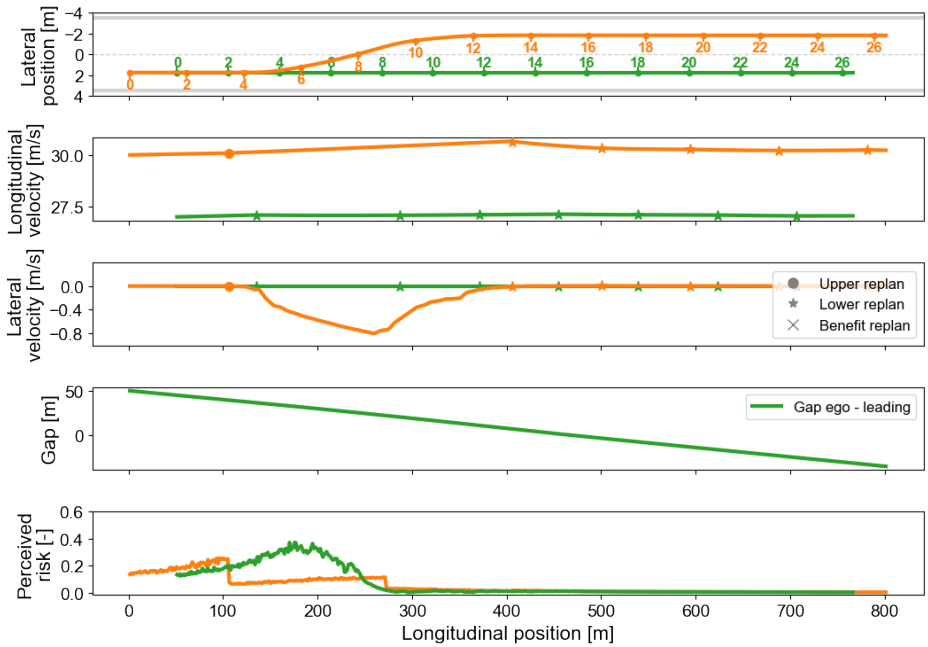

Figure B.5: Sub-scenario d). Initial conditions: the ego vehicle starts at $v(t_0) = 30$ m/s, with a gap of 50m to the lead vehicle, which has an initial speed of $v(t_0) = 27$m/s. The moderate gap and smaller speed difference provide the ego vehicle with less time to react than sub-scenario b) but more time than sub-scenario c). After $100m$ of travel, the upper-risk threshold of the ego vehicle is exceeded and replans with a lane change.

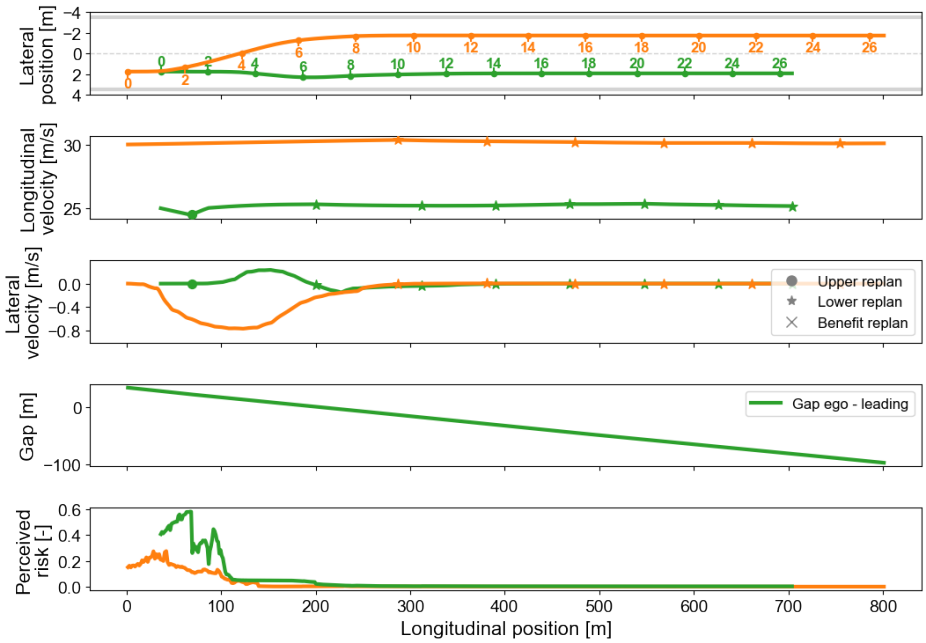

Figure B.6: Sub-scenario e). Initial conditions: the ego vehicle starts at $v(t_0) = 30$ m/s, with a gap of 30 m to the lead vehicle, which has an initial speed of $v(t_0) = 25$ m/s. Both vehicles are initialized in a high-risk scenario, characterized by close proximity and a significant relative speed. As a result, the ego vehicle immediately plans a lane change, while the lead vehicle creates more space by swerving to the right.

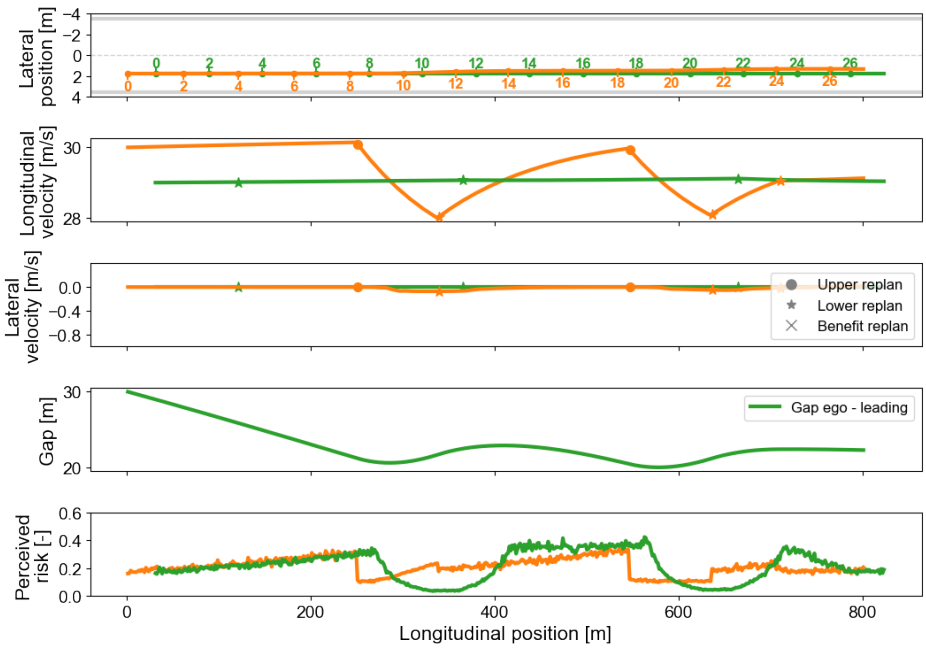

Figure B.7: Sub-scenario f). Initial conditions: The ego vehicle starts at $v(t_0) = 30$ m/s, with a gap of 30 m to the lead vehicle, which has an initial speed of $v(t_0) = 29$ m/s. The small gap and minor speed difference lead to a gradual closure, reducing the immediate incentive for the ego vehicle to overtake. As the relative speed remains low, the ego vehicle decreases its speed after surpassing the upper-risk threshold and aligns its speed with that of the lead vehicle after a lower-risk replan. Due to the risk incentive function, the risk threshold of the ego vehicle is lower, allowing the ego vehicle to plan first and avoid a collision.

B.2. Three Agents: Left-Following Vehicle Interactions

Scenario description:

This scenario demonstrates how the presence of a left-following vehicle impacts the ego vehicle's decision to overtake a slower lead vehicle, considering variations in relative distances and speeds to the left-following vehicle. In these scenarios (Figure B.8), I consider three vehicles on a two-lane road: the ego (orange) approaches a slower lead vehicle (green) on the right lane, while a left-following vehicle (blue) travels behind the ego on the left lane.

In sub-scenarios a) and b), both the ego and left-following vehicle start $50m$ apart, and the ego, in turn, is $50m$ behind the lead. The left-following vehicle's speed is either the same as the ego ($30m/s$) for a) or slightly higher ($32m/s$) for b), while the lead vehicle consistently travels at $26m/s$ across all sub-scenarios.

In sub-scenarios c) and d), the only difference is that the left-following vehicle is now only $20m$ behind the ego instead of $50m$. Finally, in sub-scenarios e) and f), the ego and left-following start at the same initial position at different lanes, with a distance of $50m$ to the leading vehicle. Again, the left-following vehicle's speed alternates between matching the ego's $30m/s$ or exceeding it at $32m/s$.

Across these configurations, the ego's decision to overtake the slower lead depends on whether there is sufficient space in the left lane, mainly if the left-following vehicle is drawing close or travelling faster, and on the benefit of passing a slower lead (at $26m/s$).

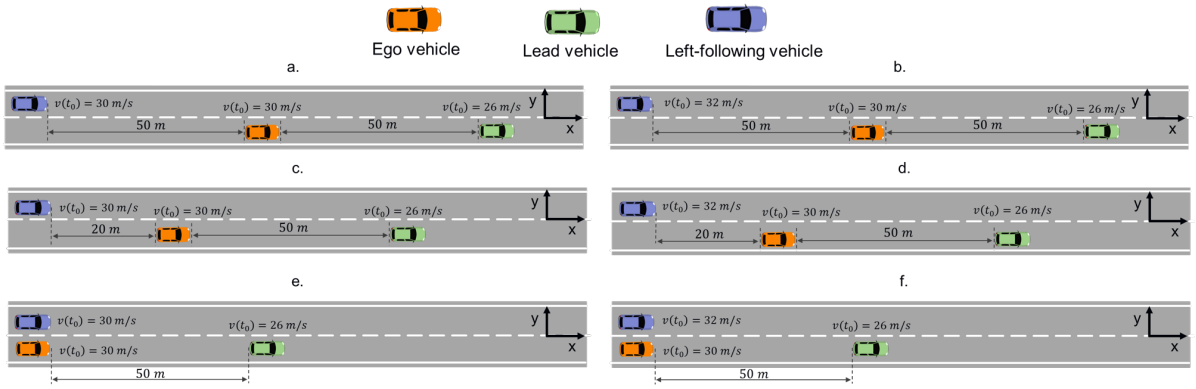

Figure B.8: Three-vehicle scenarios on a two-lane road involving an orange ego vehicle following a slower green lead vehicle and a blue left-following vehicle on the adjacent left lane. The scenarios examine variations in the left-following vehicle's distance to the ego ($50m$, $20m$, or $0m$) and its speed ($30m/s$ or $32m/s$), along with the ego's interaction with the slower lead vehicle.

Expected behaviour:

In sub-scenarios a) and b), where the ego and the left-following vehicle each start $50m$ from one another, it is expected that the ego can change lanes without cutting off the left-following vehicle. In these cases, the left-following vehicle either travels at the same speed ($30m/s$) or slightly faster ($32m/s$). Hence, the ego vehicle is expected to change in front of the left-following vehicle if there is enough time to complete the manoeuvre before the left-following vehicle closes the gap.

In scenarios c) and d), the left-following vehicle starts only $20m$ behind the ego, making it more difficult to merge in front without potentially creating a risky gap. Suppose the ego still wishes to pass the slower lead vehicle ($26m/s$). In that case, it must decide whether to accelerate sufficiently to enter the left lane ahead of the left-following vehicle or to wait until the left-following vehicle has passed and change lanes behind.

Similarly, in sub-scenarios e) and f), the ego and left-following vehicle start at the same longitudinal distance. In these cases, the ego vehicle is expected to change lanes behind the left-following vehicle. Essentially, as the left-following closes more rapidly or the initial spacing decreases, it is expected that the ego finds it riskier to merge in front and may delay a lane change or accept remaining behind the lead until the adjacent gap is more favourable. Hence, it is expected that a delayed lane change leads to a better situation for the ego vehicle.

Observed behaviour:

Across all six sub-scenarios (Figure B.9-B.14), the observed outcomes align with the expected behaviour of lane-changing benefits and risks when a third vehicle is present.

In sub-scenarios a) (Figure B.9) and b) (Figure B.10), where the left-following vehicle starts $50m$ behind, the ego completes a lane change manoeuvre without undue conflict. However, in b), the slightly faster left-following vehicle forces a speed adjustment when the ego merges in front. By contrast, in sub-scenarios c) (Figure B.11) and d) (Figure B.12), a reduced initial gap of $20m$ makes changing lanes ahead of the left-following vehicle riskier, prompting the ego to slow and delay its lane change until the left-following vehicle passes.

A similar strategy emerges in sub-scenarios e) (Figure B.13) and f) (Figure B.14), where the ego and left-following vehicle start side by side with only $50m$ to the slower lead: the ego initially remains behind the left-following vehicle, waiting for the moment to overtake the lead. In f), the higher speed of the left-following vehicle ($32m/s$) further eases the ego's decision to merge in behind.

These results confirm that larger gaps or matched speeds (sub-scenarios a and b) allow merges in front. In contrast, closer spacing or faster left-following speeds (sub-scenarios c, d, e, and f) drive the ego to merge behind or temporarily reduce its speed. Overall, the ego's decisions reflect the expected balance between overtaking the slower lead and avoiding conflicts with the adjacent lane, demonstrating that the model consistently responds to the distance and relative speed of surrounding traffic.

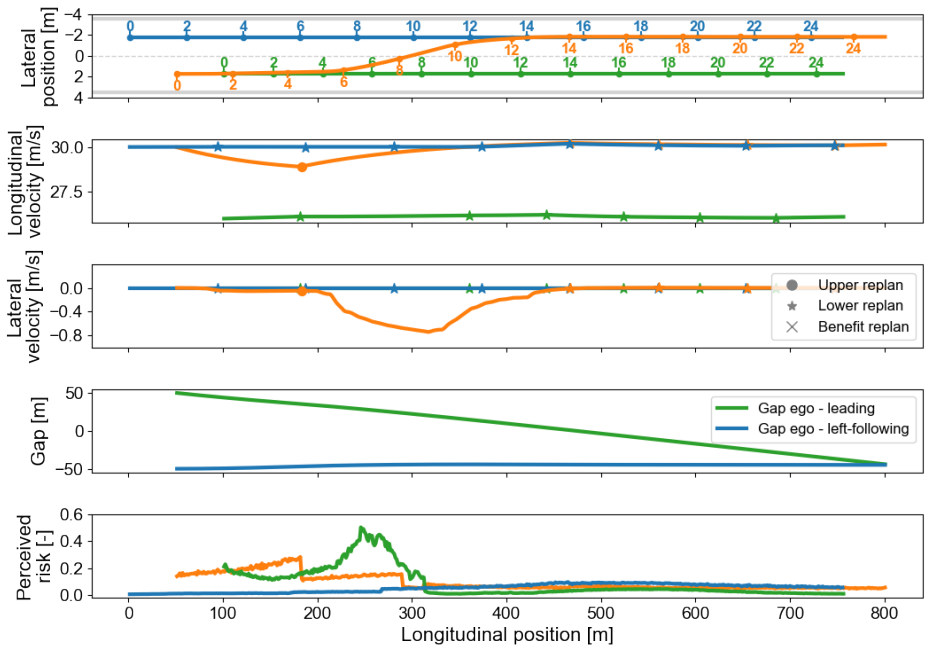

Figure B.9: Sub-scenario a). Initial conditions: The ego vehicle starts at $v(t_0) = 30$m/s, with a gap of 50m to the lead vehicle, which has an initial speed of $v(t_0) = 26$m/s. The left-following vehicle starts 50m behind the ego at $v(t_0) = 30$m/s. The ego initially reduces its speed slightly, then exceeds the upper-risk threshold and successfully performs a lane change, maintaining sufficient space ahead of the left-following vehicle.

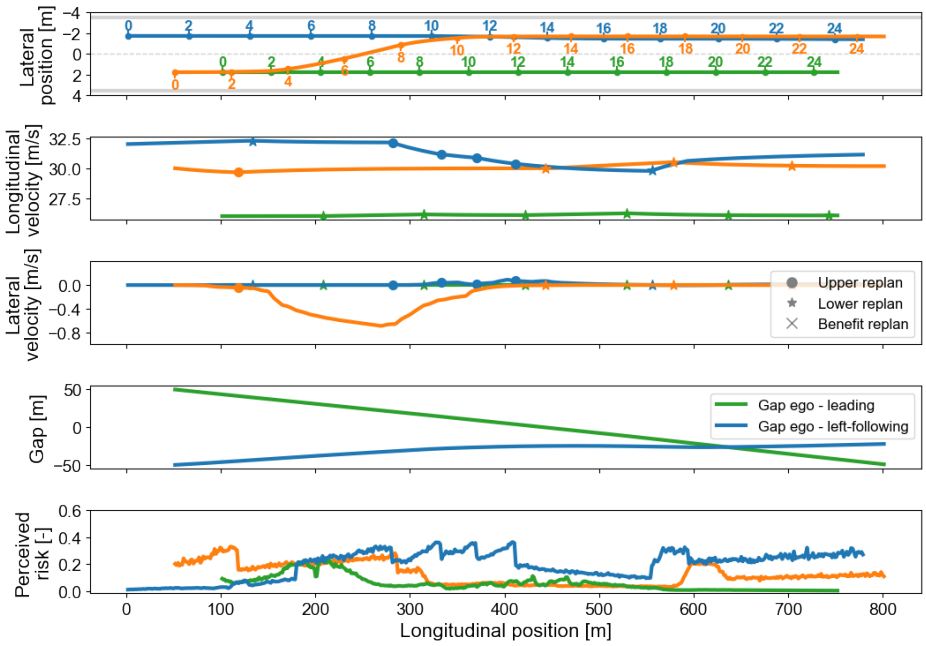

Figure B.10: Sub-scenario b). Initial conditions: The ego vehicle starts at $v(t_0) = 30$m/s, with a gap of 50m to the lead vehicle, which has an initial speed of $v(t_0) = 26$m/s. The left-following vehicle starts 50m behind the ego at $v(t_0) = 32$m/s. The ego initiates a lane change but fails to accelerate to match the speed of the left-following vehicle, forcing the left-following vehicle to brake and align with the ego's speed.

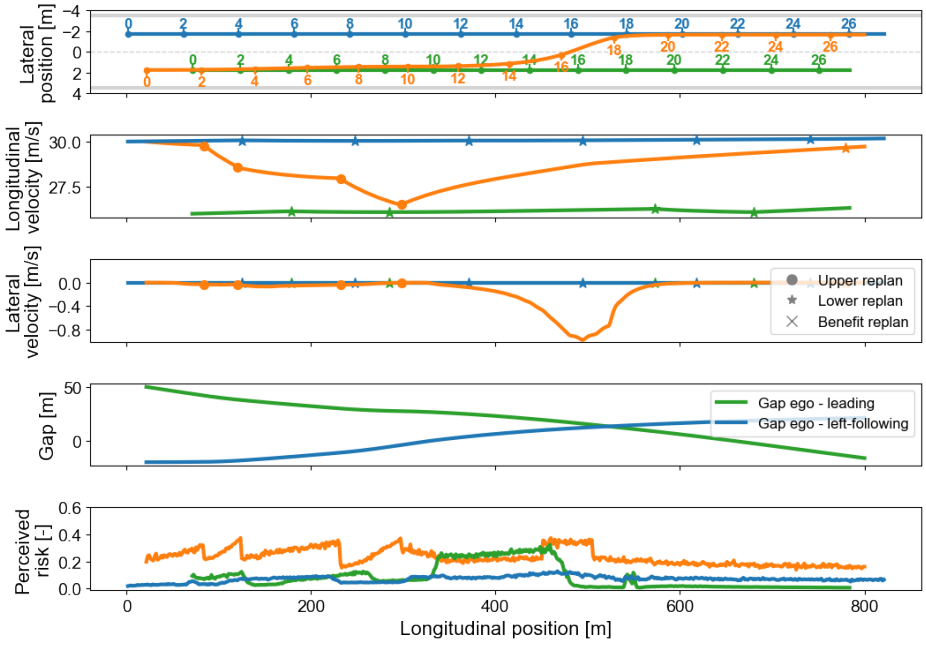

Figure B.11: Sub-scenario c). Initial conditions: The ego vehicle starts at $v(t_0) = 30$m/s, with a gap of 50m to the lead vehicle, which has an initial speed of $v(t_0) = 26$m/s. The left-following vehicle starts 20m behind the ego at $v(t_0) = 30$m/s. Due to the small gap, the ego delays its lane change, slows down to let the left-following vehicle pass, and then accelerates to its desired speed.

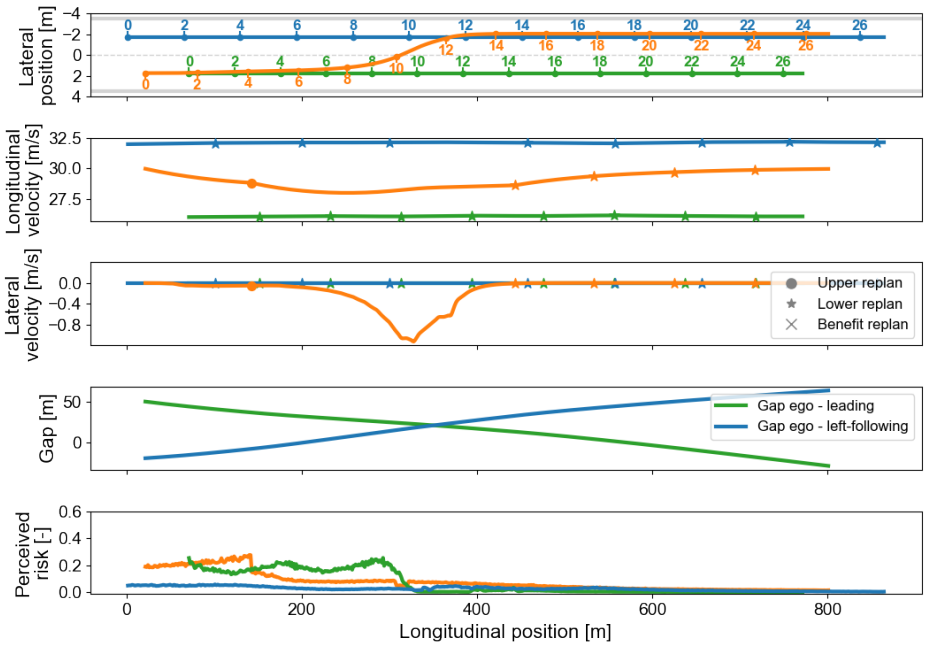

Figure B.12: Sub-scenario d). Initial conditions: The ego vehicle starts at $v(t_0) = 30$m/s, with a gap of 50m to the lead vehicle, which has an initial speed of 26m/s. The left-following vehicle starts 20m behind the ego at $v(t_0) = 32$m/s. The same scenario as in c) occurs, with the ego delaying its lane change until the left-following vehicle passes and then merging behind it. However, the higher speed of the left-following allows the ego to brake less aggressively and stay closer to its original speed before merging.

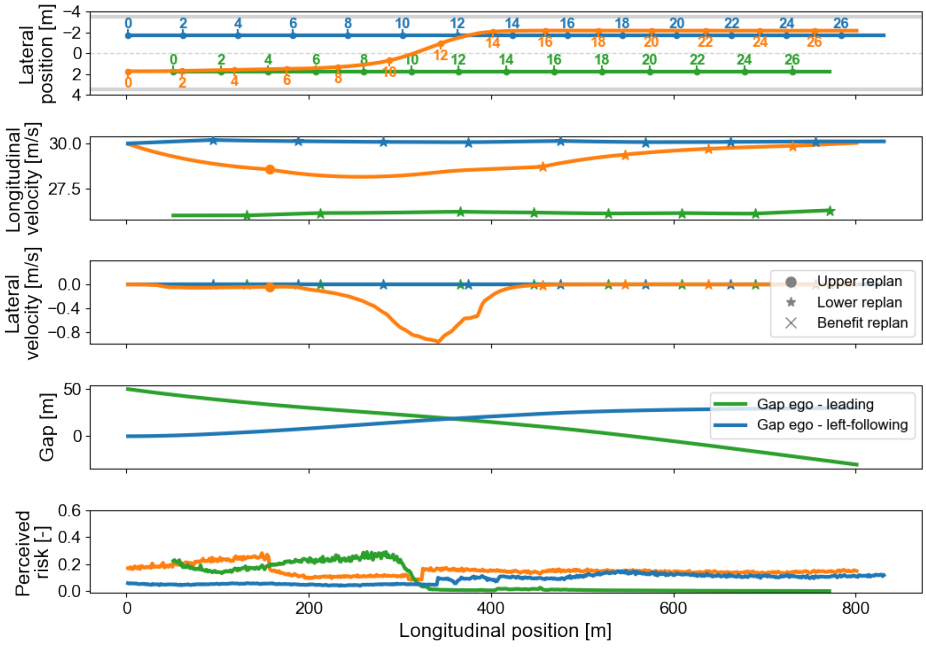

Figure B.13: Sub-scenario e). Initial conditions: The ego and left-following vehicles both start at $v(t_0) = 30$m/s, at the same longitudinal position but in different lanes, with a gap of 50m to the lead vehicle travelling at $v(t_0) = 26$m/s. The ego vehicle slightly reduces its speed, changes lanes immediately after the left-following vehicle, and then accelerates to its desired speed.

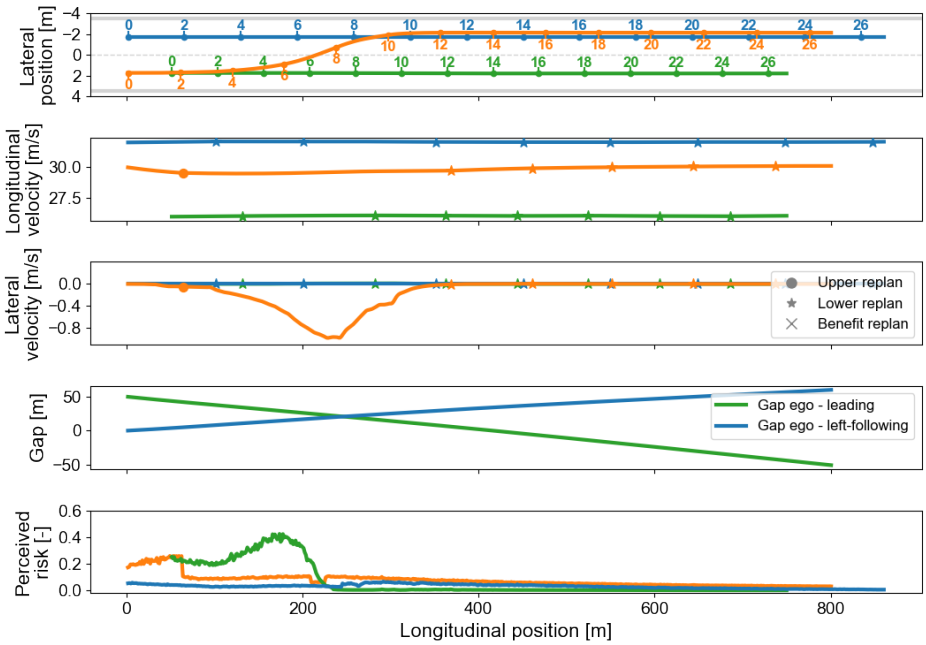

Figure B.14: Sub-scenario f). Initial conditions: The ego and left-following vehicles both start at the same longitudinal position but with different speeds: $v(t_0) = 30$m/s for the ego and $v(t_0) = 32$m/s for the left-following vehicle. The lead vehicle travels at $v(t_0) = 26$m/s with a gap of 50m to the ego. The same scenario as in e occurs; however, due to the speed difference between the ego and the left-following vehicle, the ego does not need to reduce its speed and changes lanes immediately after the left-following vehicle.

B.3. Four Agents: Complex Multi-Vehicle Interactions

Scenario description:

This scenario (Figure B.15) demonstrates the ego vehicle's behaviour in more complex traffic conditions, involving interactions with a lead vehicle (green), a left-following vehicle (blue), and a left-lead vehicle (purple), to assess how varying gaps and speeds influence lane-change strategies.

The orange ego vehicle follows a slower green lead vehicle ($26m/s$) in the right lane, initially separated by $50m$. Meanwhile, two more vehicles drive on the left lane: a purple left-lead vehicle ahead and a blue left-following vehicle behind.

In sub-scenarios a) and b), all vehicles except the green lead vehicle start with an initial speed of $30m/s$, and the distance between each vehicle is $50m$. The only difference between these two scenarios is that in scenario b), the left-following and left-leading vehicle speeds are increased to $32m/s$.

In sub-scenarios c) and d), the distance between the ego vehicle and the purple left-lead vehicle is reduced to $20m$, while all other distances remain unchanged. As in the previous cases, the left-following and left-leading vehicle alternates between $30m/s$ (scenario c) and $32m/s$ (scenario d).

Across all four sub-scenarios, the ego vehicle, travelling at $30m/s$, must decide how to overtake the green lead vehicle. This decision depends not only on whether the left lane is clear behind it (i.e., whether the left-following vehicle is present) but also on whether the left-lead vehicle is moving fast enough to make merging into the left lane a viable option.

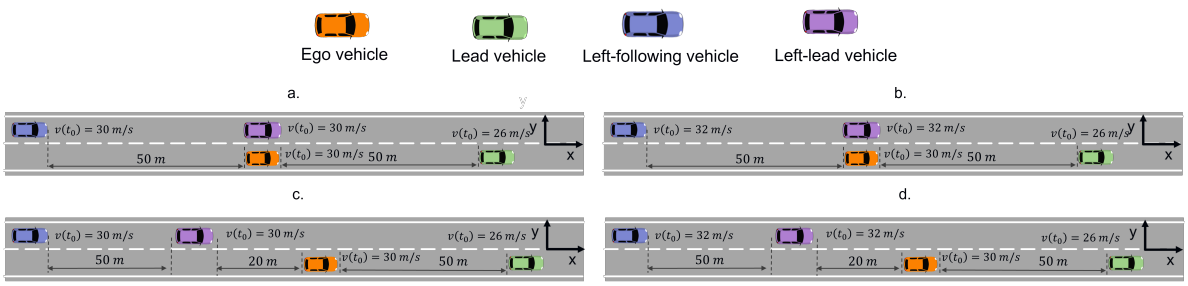

Figure B.15: Four-vehicle scenarios with an orange ego vehicle approaching a slower green lead vehicle in the right lane, a purple left-lead vehicle ahead, and a blue left-following vehicle behind in the adjacent left lane. Variations include distances between vehicles ($50m$ or $20m$) and the left-following and left-lead vehicle's speed ($30m/s$ or $32m/s$).

Expected behaviour:

It is expected that as the ego vehicle approaches the slower green lead vehicle, it will assess whether a lane change to the left is feasible. When the purple left-lead vehicle is travelling at a similar speed to the ego ($30m/s$) and the available gaps are large ($50m$), the ego should be able to merge into the left lane ahead of the blue left-following vehicle without significant disruption. However, if the blue vehicle is moving faster ($32m/s$), the ego may either accelerate quickly to merge in front or wait until the blue vehicle passes to merge safely behind.

When the gap to the purple left-lead vehicle is small ($20m$), the ego is expected to evaluate whether there is sufficient space to accelerate and avoid tailgating. A tighter available gap or a faster blue left-following vehicle should increase the risk of merging into the left lane, potentially leading the ego to delay overtaking and remain behind the green lead vehicle at $26m/s$ until a better opportunity arises.

Observed behaviour:

Across all four sub-scenarios (Figure B.16-B.19), the observed results match the anticipated balance between overtaking the slower green lead and avoiding conflicts in the left lane.

In sub-scenario a) (Figure B.16), where the blue left-following vehicle shares the same speed ($30m/s$), the ego merges left with minimal disruption. In sub-scenario b) (Figure B.17), a faster blue left-following vehicle ($32m/s$) forces a tighter merge; the ego still changes lanes but compels the left-following vehicle to reduce speed slightly.

The pattern shifts in sub-scenario c) (Figure B.18), where the purple left-lead is only $20m$ behind, and the blue left-following travels at $30m/s$. The ego stays momentarily behind the green lead until the left lane opens up, then merges between the left-leading and left-following vehicles. Finally, in sub-scenario d) (Figure B.19), the combination of a tight gap ahead ($20m$ to the purple left-lead) and a faster left-following behind ($32m/s$) prompts the ego to wait longer; once the blue vehicle passes, it merges behind both left-lane vehicles. These variations confirm that the ego's lane change strategy, whether to slip in front or merge behind, depends on the available gap to the left-lane vehicles and the relative speed of the left-lane vehicles.

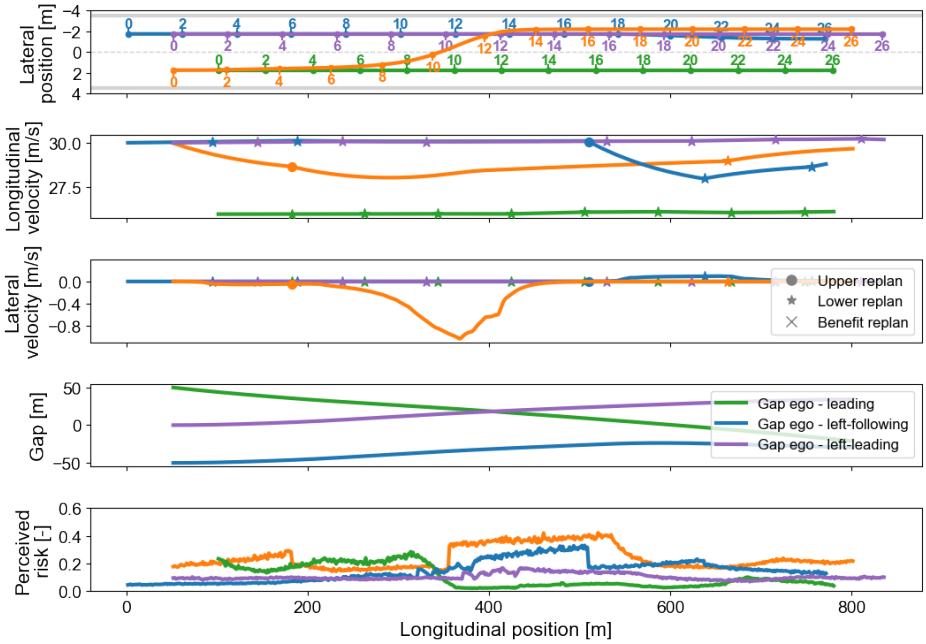

Figure B.16: Sub-scenario a). Initial conditions: The ego vehicle starts at $v(t_0) = 30$m/s, with a gap of 50m to the green lead vehicle ($v(t_0) = 26$m/s). The purple left-lead vehicle is 50m ahead of the ego, and the blue left-following vehicle is 50m behind, both at $v(t_0) = 30$m/s. The ego vehicle slightly reduces its speed and changes lanes into the left lane. It aligns behind the purple left-lead vehicle, maintaining a safe distance and avoiding disruption to the blue left-following vehicle. After completing the lane change, the ego increases its speed to continue overtaking the green lead vehicle.

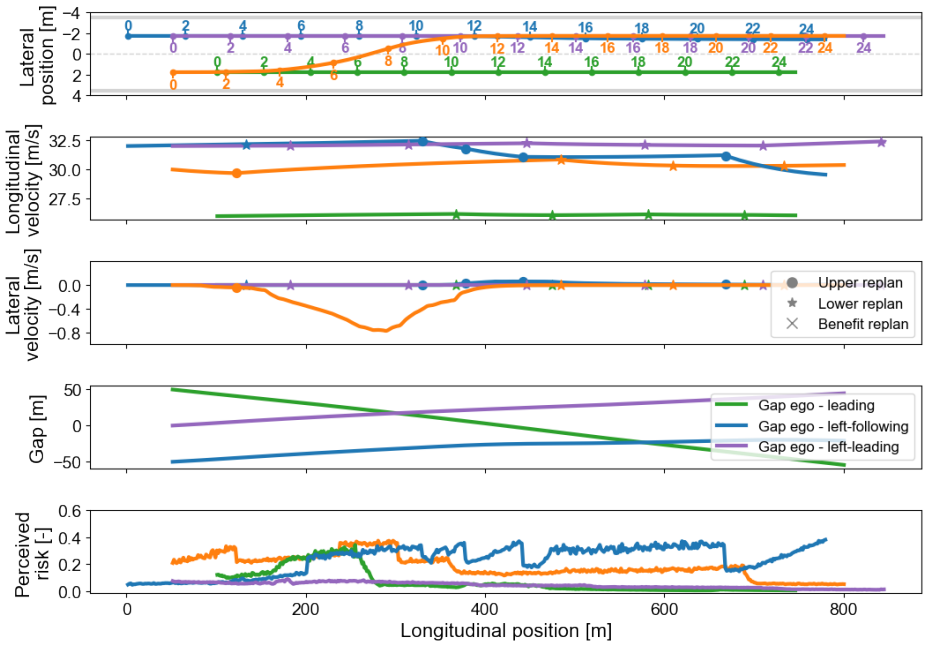

Figure B.17: Sub-scenario b). Initial conditions: The ego vehicle starts at $v(t_0) = 30$m/s, with a gap of 50m to the green lead vehicle ($v(t_0) = 26$m/s). The purple left-lead vehicle is 50m ahead, and the blue left-following vehicle is 50m behind at $v(t_0) = 32$m/s. The ego vehicle initiates a lane change in front of the blue left-following vehicle, which is travelling at a higher speed. This forces the blue vehicle to slightly reduce its speed to accommodate the ego's manoeuvre. After merging into the left lane, the ego aligns behind the purple left-lead vehicle and maintains a safe distance while overtaking the green lead vehicle.

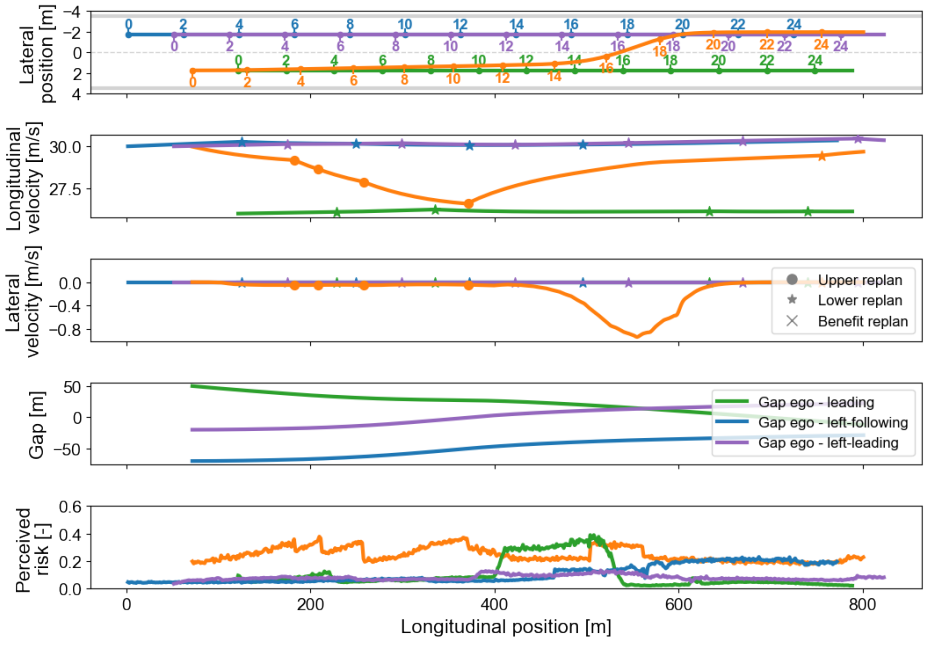

Figure B.18: Sub-scenario c). Initial conditions: The ego vehicle starts at $v(t_0) = 30$m/s, with a gap of 50m to the green lead vehicle ($v(t_0) = 26$m/s). The purple left-lead vehicle is 20m ahead, and the blue left-following vehicle is 50m behind at $v(t_0) = 30$m/s. The ego vehicle reduces its speed and stays behind the green lead vehicle in the right lane. It delays the lane change until the spacing in the left lane becomes more favourable. Once the purple left-lead has passed, the ego initiates a lane change in front of the blue left-following vehicle and accelerates to overtake the green lead vehicle.

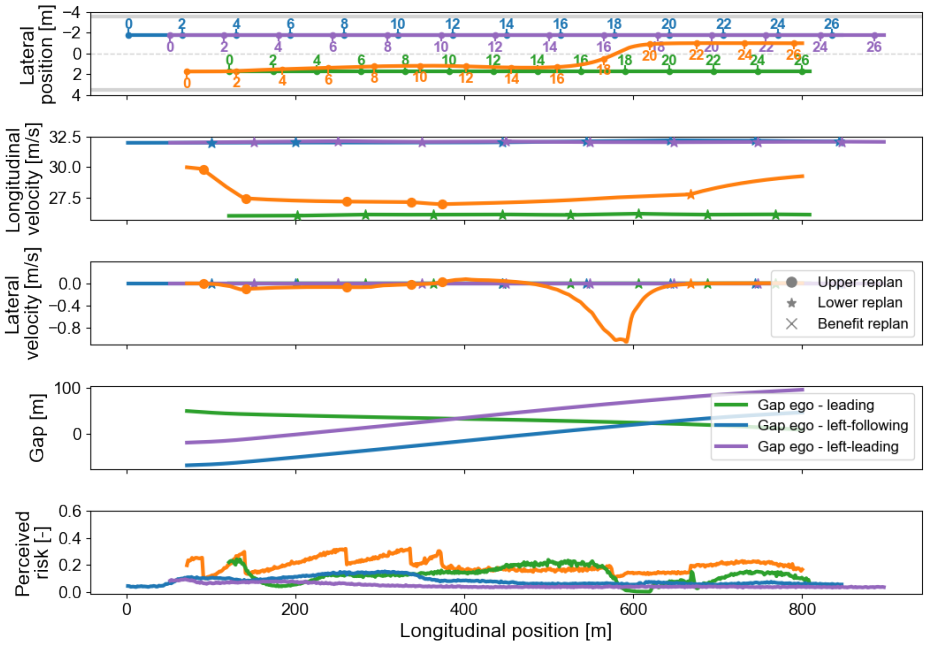

Figure B.19: Sub-scenario d). Initial conditions: The ego vehicle starts at $v(t_0) = 30$m/s, with a gap of 50m to the green lead vehicle ($v(t_0) = 26$m/s). The purple left-lead vehicle is 20m ahead, and the blue left-following vehicle is 50m behind at $v(t_0) = 32$m/s. The ego delays its lane change due to the faster blue left-following vehicle and the tight gap to the purple left-lead vehicle. After the blue vehicle passes, the ego merges into the left lane behind both the blue left-following vehicle and the purple left-lead vehicle. The ego maintains its position behind these vehicles while continuing to follow the flow of traffic.

C

Systematic Validation Against Literature on Lane Change Behaviour

This chapter presents a systematic validation of the lane change model in controlled simulations against literature on lane change behaviour. Research on discretionary lane changes consistently identifies that drivers base their decisions on a combination of factors from their current and target lanes (see [1] for a review). In the current lane, drivers typically consider their own speed, the spacing of the vehicle ahead, and the difference in speed relative to that vehicle ([2]–[4]). In the target lane, common factors for the decision to change lanes are the available gap in front of and behind vehicles in the adjacent lane, as well as on the relative speeds of those vehicles [2], [5].

The primary aim of this validation is to evaluate how variations in relative speeds and the gaps between vehicles influence the motivation to change lanes and determine whether a lane change is executed with the presented model. I investigated two categories of factors:

- Factors related to the current lane: How the speed difference and gap between the ego vehicle and a slower vehicle ahead affect the decision to initiate a lane change.
- Factors related to the target lane: How the size of the gap and the speed difference with a vehicle following in the adjacent lane influence the accepted gap in front or behind that vehicle.

Following the literature, these validations focus on relative speeds and spatial gaps as the main determinants of lane change decision-making behaviour. Although other factors, such as traffic density and vehicle size, are recognised as important in lane change motivation [1], [3], [5], this validation is limited to three agents: the ego vehicle, a lead vehicle, and a left following vehicle. This setup isolates the effects of these specific interactions without additional traffic influences. As a result, the analysis only considers the relative speed and distance to the lead vehicle in the current lane and the left following vehicle in the target lane.

By comparing the model outcomes with the expected behaviour reported in the literature, this validation demonstrates that the model reproduces the aspects of driver decision-making related to relative speeds and gaps. All simulations follow the parameter settings of Appendix A. Each scenario is run 20 times to capture the variability introduced by the stochastic belief construction and uniformly distributed parameters.

C.1. Effect of Varying Lead Vehicle Gap and Speed

Scenario Description:

The goal of this scenario (Figure C.1) is to assess how the ego vehicle's lane-change decisions are influenced by the initial gap (Δx_1) and speed difference (Δv_1) relative to a lead vehicle while keeping external influences minimal. By systematically varying Δx_1 ($30, 40, 50$, or $60m$) and Δv_1 (0 to $4m/s$), the simulation isolates the impact of these two factors on the ego's timing and likelihood of lane-changing.

The scenario consists of an ego vehicle (orange) travelling behind a slower lead vehicle (green) in the same lane on a two-lane highway. A left-following vehicle (blue) is positioned $\Delta x_2 = 50m$ behind the ego in the adjacent lane and travelling at the same speed as the ego vehicle($\Delta v_2 = 0m/s$). This ensures that the left-following vehicle never "catches up" to the ego, reducing its influence on the ego's decision to change lanes. The leading vehicle in each scenario drives at $26m/s$, and the ego and left-following vehicle speeds are adjusted based on the relative speed Δv_1.

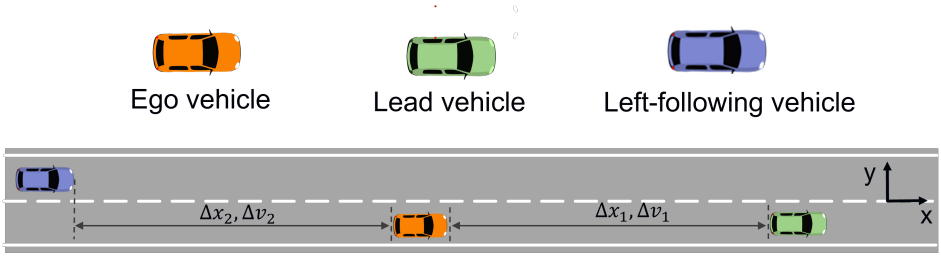

Figure C.1: Overview of the simulation setup designed to validate the influence of the initial gap (Δx_1) and speed difference (Δv_1) between the ego vehicle (orange) and the lead vehicle (blue) on lane-change decisions.

Expected Behaviour:

Empirical studies have shown that even moderate increases in the speed difference can lead to a shift in driver behaviour, with drivers opting to change lanes to gain speed advantages [1], [2]. When the ego vehicle has little or no speed advantage over the lead (e.g., $\Delta v_1 \approx 0$), the incentive to overtake is minimal, as there is no significant speed benefit. For small speed differences (approximately 0 to 2 m/s), a lane change offers limited advantages, and the ego vehicle is expected to stay in lane, matching the lead vehicle's speed and maintaining a safe following distance. However, as Δv_1 increases to 2–4 m/s, the ego vehicle becomes more motivated to overtake, given the greater speed discrepancy. In this case, it is expected that the ego vehicle changes lanes more often.

Finally, studies of naturalistic driving also highlight variability in how drivers interpret and act upon these factors [1], [5]. Differences in individual risk thresholds in the decision-making process lead to heterogeneity in the timing and likelihood of lane changes. Consequently, while larger speed differences generally increase the frequency of lane changes, some drivers will remain in the current lane under conditions where others would have already initiated an overtaking manoeuvre.

Observed Behaviour:

Figure C.2 shows how the fraction of simulations, including lane changes, increases with Δv_1. As expected, at or near $0m/s$ difference, no lane changes occur since the vehicles travel at nearly the same speed. At moderate Δv_1 values ($1-2m/s$), the lane-change rate rises only for vehicles initialized close to the lead vehicle ($\Delta x_1 = 30$ m). At higher Δv_1 values ($3-4m/s$), nearly all runs involve a lane change.

Still, the repeated trials reveal variability in the ego vehicle's decision-making: in some instances, the ego remains in its current lane to avoid unnecessary manoeuvres, while in other simulations, it overtakes more decisively, reflecting subtle differences in risk thresholds and the stochastic belief construction. This effect can be observed when $\Delta x_1 = 30$ m and Δv_1 is small (around 1 m/s), where a subset of trials shows a lane change, and others do not.

Figure C.3 presents the time headway to the lead vehicle at the moment when a lane change starts as a function of Δv_1. Several observations stand out. Larger Δv_1 often coincides with a higher time headway at the manoeuvre start. Because the ego recognizes a strong incentive to overtake, it plans the manoeuvre before the headway grows dangerously small. Thus, although the lane change may occur sooner in absolute time, it still happens at a more comfortable, i.e., larger, headway. Additionally, there is greater variability at higher Δx_1. For initial distances of $50m$ or $60m$, the ego can either wait or pass relatively early, resulting in higher time headway values. There is less time to react with shorter initial gaps, such as $30m$, so most drivers initiate a swift pass. As a result, lane changes at this distance are planned almost immediately, leading to values that are closely together.

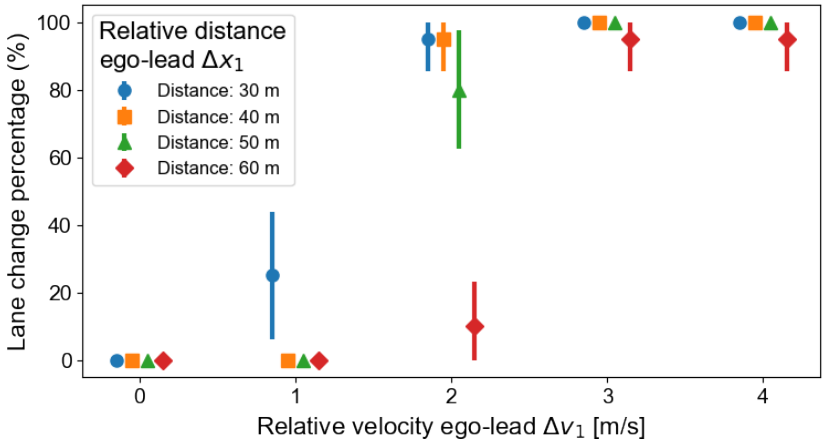

Figure C.2: Fraction of simulations that result in a lane change as a function of the initial speed difference Δv_1 to the lead vehicle, grouped by initial gap Δx_1. Error bars denote binomial proportion standard error of mean.

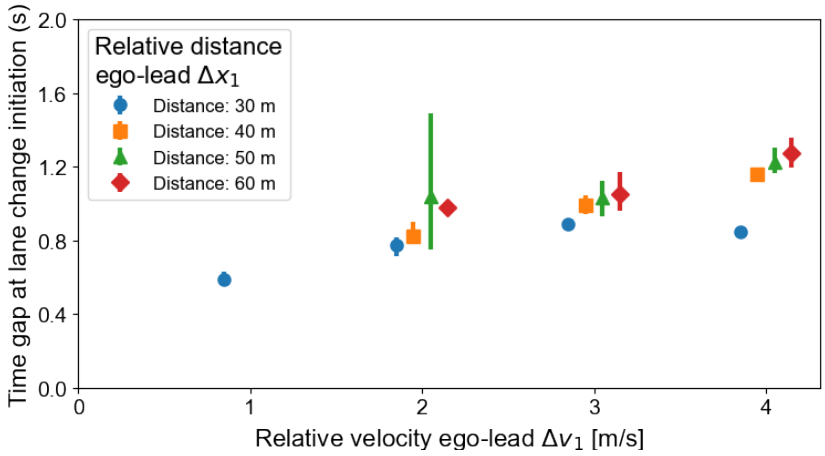

Figure C.3: Time headway to the lead vehicle at the start of a lane change versus Δv_1. The error bars represent interquartile ranges.

C.2. Effect of Varying Left-Following Vehicle Gap and Speed

Scenario Description:

The goal of this scenario is to assess how the gap (Δx_2) and speed difference (Δv_2) of a left-following vehicle influence the ego vehicle's decision to change lanes, particularly whether ego accepts the gap ahead of or behind the left-following vehicle.

The simulation setup is similar to the one described in Section C.1, where an ego vehicle (orange) follows a lead vehicle (green) in the same lane on a two-lane highway. A left-following vehicle (blue) is included in the left adjacent lane, as illustrated in Figure C.1. However, unlike the previous setup, the lead vehicle is positioned at a fixed distance of $\Delta x_1 = 60m$ ahead of the ego vehicle and travels $4m/s$ slower than the ego vehicle ($\Delta v_1 = 4m/s$). The ego vehicle maintains a constant initial velocity of $30m/s$, ensuring that a lane change is always prompted due to the velocity difference between the ego and the lead vehicle.

The influence of the left-following vehicle, travelling in the adjacent lane, is studied by systematically varying the relative distance Δx_2 with values of $0, 20, 40,$ and $60m$, and the relative velocity Δv_2, which ranges from -4 to $2m/s$, meaning the left-following vehicle can be up to $4m/s$ faster or $2m/s$ slower. This configuration isolates how the gap and speed of the left-following vehicle impact the ego vehicle's lane change decision. All simulations follow the parameter settings of Appendix A.

Expected Behaviour: Empirical research suggests that gap acceptance during lane changes depends on both the spatial gap (Δx_2) and the relative speed difference (Δv_2) between the ego vehicle and the left-following vehicle [2], [5].

When the approaching vehicle in the target lane is moving substantially faster (e.g., $\Delta v_2 \ll 0$), drivers typically delay the lane change until the faster vehicle passes, thereby reducing potential conflicts or rear-end collision risks. Conversely, when the relative speed is similar or slightly in the ego vehicle's favour (i.e., $\Delta v_2 \approx 0$ or positive), drivers may be more inclined to merge ahead, provided there is sufficient space to accelerate and maintain a safe following distance once in the new lane [5].

Observed Behaviour: Figure C.4 shows how the fraction of simulations, including a lane change, is performed either in front of or behind the left-following vehicle. As expected, the number of lane changes performed in front increases with Δv_2. For large negative Δv_2 (e.g., $-4m/s$), lane changes are most frequent behind the left-following vehicle when $\Delta x_2 = 40m$ or $60m$. This suggests that faster-moving left-following vehicles create favourable conditions for the ego vehicle to merge behind them. For smaller Δv_2 values (e.g., -1 m/s to 0 m/s), lane changes are increasingly performed ahead of the left-following vehicle, particularly at larger Δx_2. This reflects the ego vehicle's ability to change into the adjacent lane without interfering with the left-following vehicle's trajectory.

In addition to these aggregate trends, the simulation data exhibit variability between each trial, highlighting how differences in risk thresholds can alter the outcome. Even under similar initial conditions (e.g., a moderate negative Δv_2 and medium spacing), some ego vehicles will opt to merge behind, while others will choose to merge in front if they judge the gap to be sufficient. This variability is most evident at transitional values of Δv_2, where neither merging in front nor merging behind is preferred across all trials. As the relative velocity approaches positive or near-zero values, the ego vehicle's ability to comfortably merge in front becomes more appealing, provided the spacing is sufficient. Conversely, when the left-following vehicle is much faster (large negative Δv_2), merging behind offers a lower-risk option.

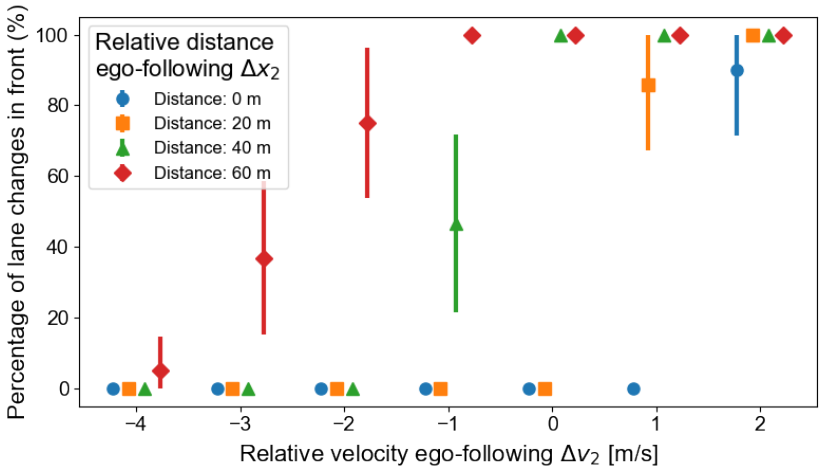

Figure C.4: Fraction of simulations that result in a lane change in front of the left-following vehicle as a function of the initial speed difference Δv_2 to the left-following vehicle, grouped by initial gap Δx_2. Error bars denote binomial proportion standard error of mean.

D

Grid Search

In order to reproduce observed tactical behaviours from the highD dataset, two parameters of the benefit-based trigger, the drift rate α and the threshold θ_s, were calibrated via a grid search. By tuning these parameters, I aimed to identify the settings under which simulated driver decisions most closely resemble the real-world variability in lane-change decision and their timing.

D.1. Scenario and Dataset Extraction

A representative scenario from the highD dataset was selected to illustrate the decision-making process and to serve as the basis for the grid search. In this scenario (see Figure D.1), the ego vehicle is approaching a slower-moving truck. This situation frequently motivates an early lane change: the ego driver anticipates slower traffic ahead and decides to switch lanes in advance to avoid being impeded by the truck.

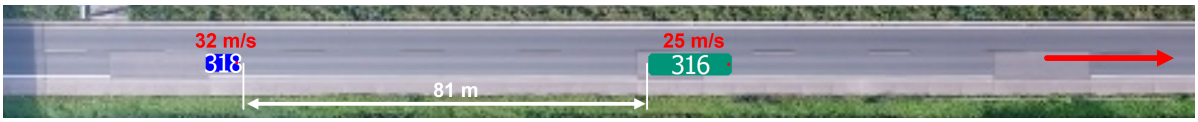

Figure D.1: Traffic scenario (highD dataset 20, frame 6103, ego vehicle id 318) used to extract similar samples using the Hausdorff extraction method. The blue vehicle is the ego vehicle, with the red arrow indicating the driving direction. The green vehicle denotes a slow-moving truck, all driving in the same direction on the two-lane highway.

To obtain variability in tactical behaviours under similar initial conditions, I applied the Hausdorff extraction method to the highD dataset [6]. This method identifies samples from the dataset most similar to the representative scenario in terms of vehicle positions and velocities. As a result, 500 samples were extracted, each reflecting comparable traffic conditions with variability in driver responses (Figure D.2).

Among these 500 extracted samples, the obtained tactical variability contains 136 samples containing a lane change, and 164 samples remain in the initial lane (see Figure D.3). This distribution of responses captures realistic, human-like decision-making in the presence of a slower-moving truck.

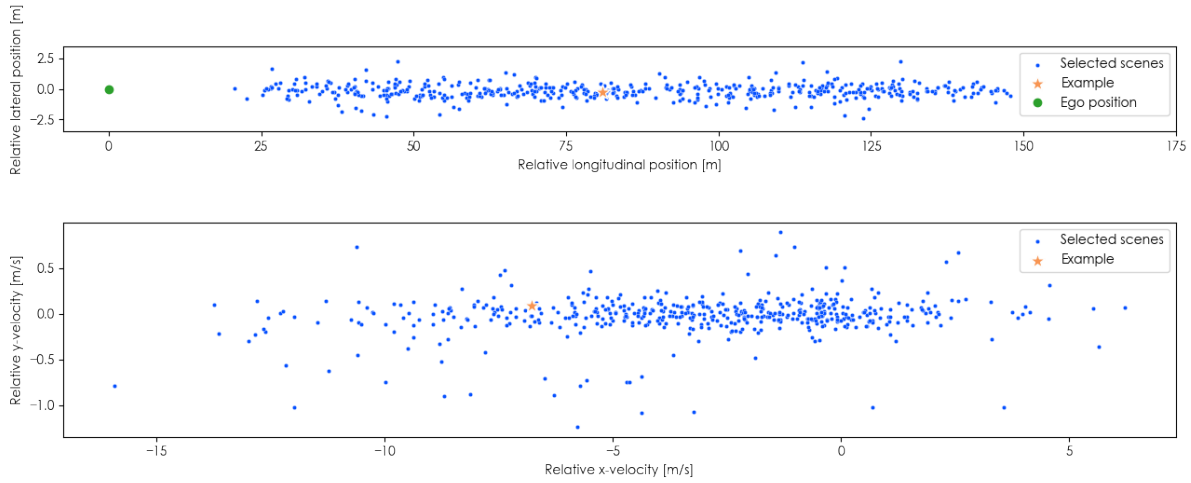

Figure D.2: The spread of the blue points represents the samples obtained after the Hausdorff extraction. The top plot shows the positions of surrounding vehicles relative to the ego vehicle, where the starts represent the scenario of interest. The bottom plot shows the relative velocities of the surrounding vehicles (denoted by the stars) to the ego vehicle. The blue dots represent the 500 closest context sets automatically extracted from the highD dataset.

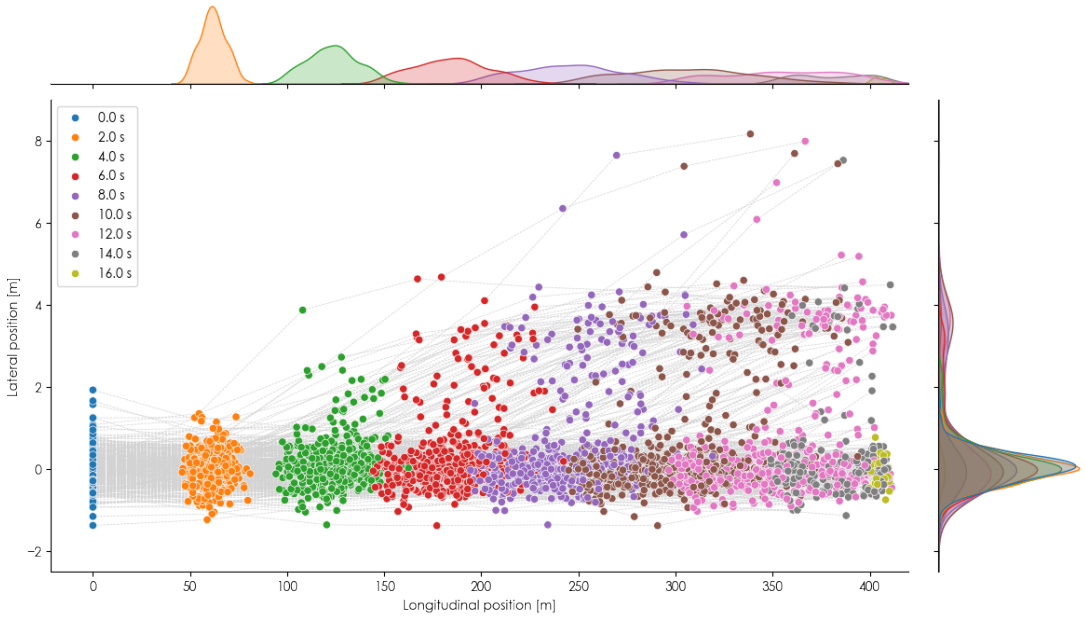

Figure D.3: The variability in driver responses as they evolve from the 500 extracted samples.

D.2. Grid Search Setup and Parameter Ranges

To systematically determine the influence of α and θ_s, I performed a grid search over a predefined set of values. I used a 10×10 grid for every condition using the drift rate α in the range [0.005, 0.05] and the threshold θ_s in the range [0.5, 5.0]. The boundary B was held constant at $B = 1$ to reduce computational complexity, the numbers of simulations, and to isolate the effects of changing α and θ_s.

Thus, for each of the 10 α-values and 10 θ_s-values, a total of $10 \times 10 = 100$ parameter configurations were evaluated. Under each configuration, 500 simulations were run using the extracted samples (for a total of 50,000 simulations). In these simulations, all other parameters remained fixed, and lane changes that risk considerations might trigger were limited by lowering the weight penalty on deviations from desired velocity ($\lambda_1 = 0.1$). This ensures that the observed lane-change decisions are driven mainly by the benefit lane-change trigger being investigated.

D.3. Performance Metrics

To identify the best-fitting parameters, I evaluated each configuration against two performance metrics:

- **F1-score:** This metric measures how well the simulated lane-change decisions (lane-change vs. no lane-change) match the responses observed in the dataset. It combines precision (the fraction of correct positive predictions) and recall (the fraction of actual positive occurrences correctly identified).
- **Time headway at initiation:** This refers to the time headway between the ego vehicle and the slower-moving lead vehicle when a lane change is initiated.
- **Combined metric:** This metric integrates the F1-score and the time headway error into a single performance indicator. It is calculated using the equation:

$$\text{Combined Metric} = \frac{\text{F1-score}}{1 + \frac{1}{\max(\text{THW error})} \times \text{THW error}}. \tag{D.1}$$

D.4. Results

Based on the grid search results, the best-performing parameter range for the drift rate α lies between 0.02 and 0.04, and for the threshold θ_s, the range is between 3.0 and 4.0, as indicated by the high combined metric values (blue regions) in the rightmost plot (Figure D.4). While the results are not entirely convex, these regions exhibit the best performances.

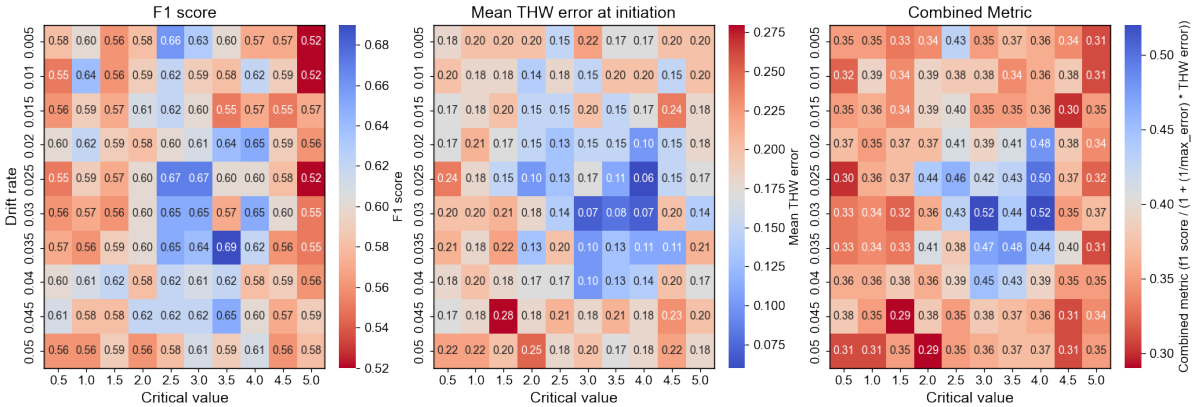

Figure D.4: Results of the grid search on the drift rate α and threshold θ_s. The f1 score heatmap (left) shows how the simulated lane change decisions align with the dataset's observed responses. The mean time headway error heatmap (centre) evaluates the timing/initiation of lane changes. The combined metric (right) integrates both metrics into a single performance indicator. Higher values (blue regions) indicate better performance across all metrics.

E

Validation Without Benefit-Based Trigger

When the benefit-based trigger is disabled, the model relies solely on risk-based replanning (by exceeding the upper-risk threshold) to initiate lane changes. As shown in Table E.1, this approach reduces the overall number of lane changes in scenario A from 36.2% to 29.0%. In Scenario B, it likewise lowers the lane-change rate from 22.4% to 14.0%, bringing it closer to but not exactly matching the human rate in that scenario (15%).

Figure E.1 presents the time-gap distributions at the moment of lane-change initiation for the lead, left-following, and left-leading vehicles, aggregating data from both scenarios. Only samples in which both the model and human drivers chose to change lanes are included. Each subplot indicates how the model's distribution (without the benefit trigger) shifts relative to the human distributions. Notably, the median time gaps for the ego–lead and ego–left-following comparisons are smaller in the model. For instance, the median ego–lead gap was $1.1s$ when the benefit-based trigger was active, dropping to $0.95s$ without it.

These smaller time gaps suggest a possible reason for the reduction in lane changes: the model does not initiate a lane change when the lead vehicle is still relatively far away, whereas the benefit-based trigger would have prompted an earlier lane change. Furthermore, the executed lane changes tend to be initiated at smaller time gaps. As a result, relying solely on risk-based replanning often leads to later or fewer lane changes than humans typically perform. Thus, incorporating both risk- and benefit-based triggers allows the model to better reproduce human behaviour.

Table E.1: Tactical behaviour of the ego vehicle in scenarios A (high-speed-difference) and B (low-speed-difference), comparing lane-change, lane-keeping, off-road events, and collision rates between the model with benefit-based replanning, ablation without, and human drivers.

	Scenario A: high-speed-difference			**Scenario B: low-speed-difference**		
	Percentage of model behaviour	Percentage without benefit trigger	Percentage of human behaviour	Percentage of model behaviour	Percentage without benefit trigger	Percentage of human behaviour
Lane-change	36.2 ± 1.2	29.0	36.0	22.4 ± 2.7	14.0	15.0
Lane-keeping	63.7 ± 1.2	71.0	64.0	77.5 ± 2.8	86.0	85.0
Off-road	0.1 ± 0.2	0.0	0.0	0.1 ± 0.1	0.0	0.0
Collision	0.0 ± 0.0	0.0	0.0	0.0 ± 0.0	0.0	0.0

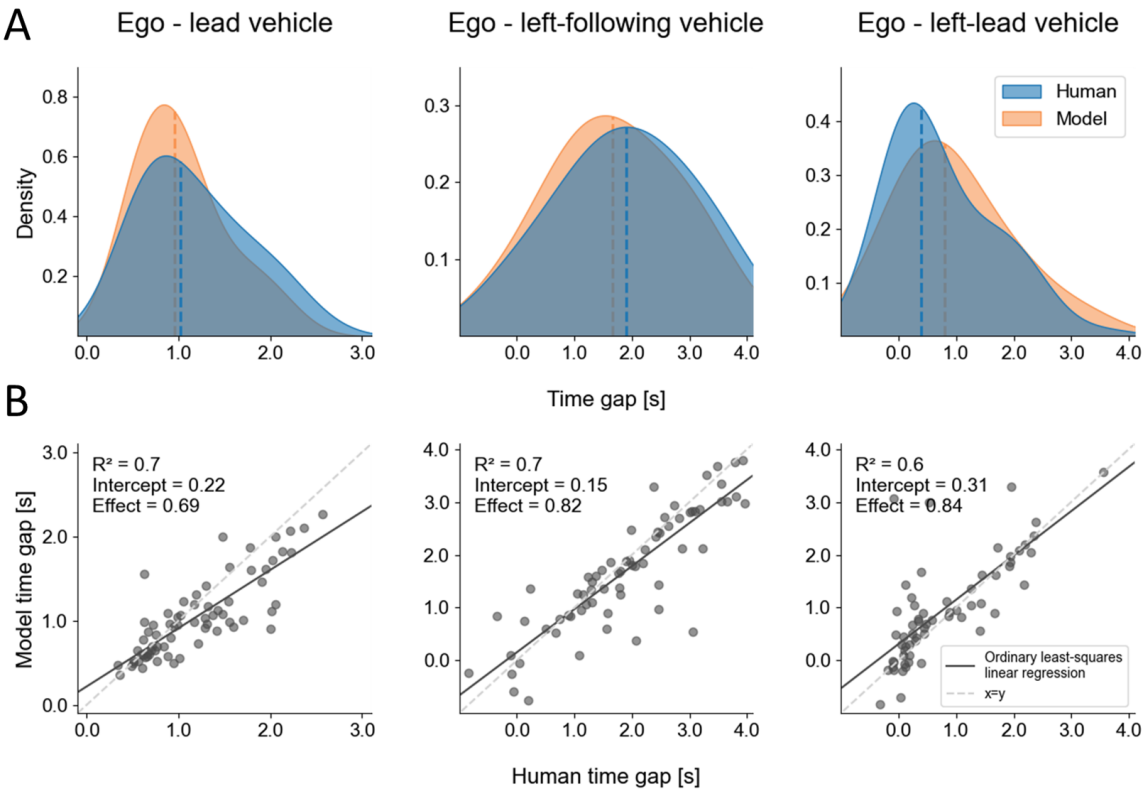

Figure E.1: An overview of the accepted gaps at lane change initiation, without the benefit-based trigger. **A)** Estimated distributions of time gap at the moment of lane change initiation to the lead, left-following, and left-leading vehicle. Dashed lines indicate the medians of each distribution. **B)** The relationships between the human and model behaviour for all data points corresponding to the time gap distributions.

Bibliography

[1] Y. Ali, Z. Zheng, M. Mazharul Haque, *et al.*, "Understanding the discretionary lane-changing behaviour in the connected environment," en, *Accident Analysis & Prevention*, vol. 137, p. 105 463, Mar. 2020.

[2] S. Moridpour, G. Rose, and M. Sarvi, "Effect of Surrounding Traffic Characteristics on Lane Changing Behavior," en, *Journal of Transportation Engineering*, vol. 136, no. 11, pp. 973–985, Nov. 2010.

[3] M. Guo, Z. Wu, and H. Zhu, "Empirical study of lane-changing behavior on three Chinese freeways," en, *PLOS ONE*, vol. 13, no. 1, X. Ma, Ed., e0191466, Jan. 2018.

[4] E. Balal, R. L. Cheu, T. Gyan-Sarkodie, *et al.*, "Analysis of Discretionary Lane Changing Parameters on Freeways," en, *International Journal of Transportation Science and Technology*, vol. 3, no. 3, pp. 277–296, Sep. 2014.

[5] M. Yang, X. Wang, and M. Quddus, "Examining lane change gap acceptance, duration and impact using naturalistic driving data," en, *Transportation Research Part C: Emerging Technologies*, vol. 104, pp. 317–331, Jul. 2019.

[6] O. Siebinga, A. Zgonnikov, and D. Abbink, "Uncovering Variability in Human Driving Behavior Through Automatic Extraction of Similar Traffic Scenes from Large Naturalistic Datasets," en, in *2023 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, Honolulu, Oahu, HI, USA: IEEE, Oct. 2023, pp. 4790–4796.