

**Document Version**

Final published version

**Licence**

CC BY-NC-ND

**Citation (APA)**

Schumann, J. F., Engström, J., Johnson, L., O’Kelly, M., Messias, J., Kober, J., & Zgonnikov, A. (2026). Active inference as a model of collision avoidance behavior in human drivers. *Nature Communications*, 17(1), Article 5009. <https://doi.org/10.1038/s41467-026-73345-0>

**Important note**

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

**Copyright**

In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership. Unless copyright is transferred by contract or statute, it remains with the copyright holder.

**Sharing and reuse**

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

**Takedown policy**

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

# Active inference as a model of collision avoidance behavior in human drivers

---

Received: 2 June 2025

---

Accepted: 6 May 2026

---

Published online: 10 June 2026

---

 Check for updates

---

Julian F. Schumann<sup>1</sup>, Johan Engström<sup>2</sup>✉, Leif Johnson<sup>2</sup>, Matthew O’Kelly<sup>2</sup>, Joao Messias<sup>2</sup>, Jens Kober<sup>1</sup> & Arkady Zgonnikov<sup>1</sup>

Collision avoidance – involving a rapid threat detection and quick execution of the appropriate evasive maneuver – is a critical aspect of driving. However, existing models of human collision avoidance behavior are fragmented, focusing on specific scenarios or only describing certain aspects of the avoidance behavior, such as response times. This paper addresses these gaps by proposing a computational cognitive model of human collision avoidance behavior based on active inference. Active inference provides a unified approach to modeling human behavior: the minimization of free energy. Building on prior active inference work, our model incorporates established cognitive mechanisms such as evidence accumulation to simulate human responses in three distinct collision avoidance scenarios: front-to-rear lead vehicle braking, lateral incursion by an oncoming vehicle, and another vehicle failing to yield at an intersection. We demonstrate that our model explains a wide range of empirical findings on human collision avoidance behavior. Specifically, the model closely reproduces both aggregate results from meta-analyses previously reported in the literature and detailed, scenario-specific effects observed in two recent driving simulator studies, including response timing, maneuver selection, and execution. Our results highlight the potential of active inference as a generalizable framework for understanding and modeling human behavior in complex real-life driving tasks.

## Main

Collision avoidance is a critical skill for human drivers. It involves the rapid detection of threats (such as a vehicle ahead suddenly braking) and deciding on an appropriate evasive maneuver (for instance, braking or swerving). These maneuvers are complex, requiring not only precise execution but also continuous adjustments as the situation evolves. Furthermore, drivers need to account for the uncertainty in the future behavior of other road users: for example, will an oncoming vehicle encroaching into my lane continue across or move back to its own lane? Understanding how humans avoid collisions in traffic can provide insights into high-

stakes, split-second decision making which has substantial implications for traffic safety.

Behavior models play a key role both in understanding the mechanisms of human collision avoidance and in improving traffic safety. These models are applied in diverse contexts, such as collision risk estimation<sup>1</sup>, understanding effects of driver distraction<sup>2</sup>, modeling take-over behavior<sup>3</sup>, representing human agents in simulated test environments<sup>4</sup>, and providing behavioral benchmarks for automated vehicles<sup>5,6</sup>. Besides immediate practical applications, modeling human collision avoidance behavior is an interesting subject of study in its own right. Being a highly complex and dynamic task with extremely

---

<sup>1</sup>Department of Cognitive Robotics, Delft University of Technology, Delft, Netherlands. <sup>2</sup>Waymo LLC, Mountain View, CA, USA.

✉ e-mail: [jengstrom@waymo.com](mailto:jengstrom@waymo.com)

high stakes, collision avoidance provides a unique testbed for theories and models of cognition previously not validated in the real world<sup>7–9</sup>.

Most existing computational models of human collision avoidance are mechanistic, that is, are based on the explicit modeling of cognitive mechanisms underlying response timing and evasive maneuvering. However, they are typically fragmented, focusing either on specific scenario types (e.g., front-to-rear conflicts<sup>10–12</sup> or merging<sup>13</sup>), specific explanatory factors (such as off-road glances<sup>1,14</sup> or cognitive load<sup>2</sup>), or only reproduce certain aspects of human behavior (such as response times<sup>3,6,10,11,15</sup> or the extent of steering<sup>16</sup>). Altogether, these models cover a wide range of scenarios and diverse aspects of collision avoidance behavior. Yet, each one of these mechanistic models on its own is highly specific: they are not designed to generalize to multiple scenarios or describe multiple aspects of human behavior.

Recently, machine learning models based on large datasets of human driving have demonstrated the ability to generalize across a wide range of traffic scenarios<sup>17–22</sup>. Because such models typically generate full motion trajectories, they also have the potential to represent multiple aspects of collision avoidance behavior and not just a single metric of interest. However, a key challenge is that safety-critical behavior, such as collision avoidance, is often under-represented in the datasets used for training<sup>23</sup>, which makes it hard to achieve representative human-like collision avoidance behavior solely based on learning from data<sup>24</sup>.

Thus, there is currently a lack of models that can capture the key aspects of human collision avoidance behavior (response selection, timing, and execution) all at the same time and across different scenarios. This limits both practical applications (due to the need to develop a new model for every new scenario) and fundamental understanding of cognitive mechanisms underlying the behavior of humans in safety-critical situations in traffic (due to the lack of a unified explanation for multiple aspects of behavior).

To address this gap, here we present a model of human collision avoidance behavior based on active inference. Originating in computational neuroscience, active inference is a versatile general framework for understanding and modeling sentient behavior in living systems<sup>25–27</sup> that has been previously used to model human behavior in diverse contexts<sup>28–32</sup>, including the modeling of human driver behavior<sup>33</sup> such as car following<sup>34</sup>, responses to driving automation failures<sup>15</sup>, and managing uncertainty around occlusions and non-driving-related tasks<sup>35</sup>. Building on the model of Engström et al.<sup>35</sup>, our active inference model is designed to reproduce a wide spectrum of human behavior (e.g., chosen maneuver, reaction timing, collision likelihood, etc.) in response to potential collisions. To this end, our model incorporates several well-known cognitive mechanisms to represent the dynamics of human decision making in response to sudden stimuli, such as looming perception<sup>36,37</sup> and evidence accumulation<sup>38–41</sup>. We evaluated our model against a range of previously reported empirical findings in three paradigmatic collision avoidance scenarios: 1) the front-to-rear scenario, where a driver needs to respond to a suddenly braking vehicle in front, 2) the opposite-direction lateral incursion scenario, where an oncoming vehicle suddenly cuts across the driver's path, and 3) the intersection right-turn-into-path scenario, where a vehicle coming from the right on a perpendicular road does not yield when turning right into the driver's lane. Testing the model in three markedly different scenarios allowed us to investigate the potential of active inference as a framework for generalizable modeling of human collision avoidance behavior.

## Results

In this work, we use active inference as an overarching framework to guide the modeling. The key premise of active inference is that all behavior and cognition can be understood based on the single principle of minimizing free energy. An agent minimizing free energy can be conceptually understood as the agent sensing and acting upon the

world in such a way as to minimize its surprise over time. In the active inference framework, this amounts to seeking out observations that are expected (unsurprising) and preferred given the type of creature the agent is, reflecting its adaptation to its particular environment or niche (e.g., a fish expects and prefers the sensation of being immersed in water). Agents capable of planning into the future and modeling the consequences of their actions (such as human car drivers), select plans that minimize the expected surprise, or more generally, the *expected free energy* (EFE)<sup>25</sup>.

Our model implements these principles in the collision avoidance context (Fig. 1). Specifically, the modeled agent repeatedly evaluates possible futures under its current policy (i.e., a sequence of planned actions) – including interactions with other road users – in terms of their EFE. The driver then selects policies that minimize EFE by realizing observations that align with the driver's preferences, such as avoiding collisions (yielding *pragmatic value*) while obtaining new information to reduce uncertainty in the agent's beliefs about the future (yielding *epistemic value*)<sup>25,35</sup>.

Fundamentally, agents cannot have perfect knowledge of the mechanisms underlying the surrounding world (the *generative process*), but instead rely on an internal, not necessarily veridical, representation of their beliefs about the environment, the *generative model*<sup>25</sup>. Importantly, an (ego) agent's *generative model* is probabilistic, with the uncertainty about the state of the world being represented by a set of sample particles. Our model implements these general principles in a time-discrete, sequential process (Fig. 2) which incorporates a number of key perceptual, cognitive, and motor mechanisms.


### Looming-based perception

First, the agent observes the world and updates its belief (Fig. 2a), combining these observations with the expectations derived from the generative model propagating its previous belief forward in time. The model assumes that the agent cannot directly perceive kinematic states of surrounding objects (e.g., positions and velocities), but instead uses the readily available perceptual information<sup>37</sup>: the object's visual angle  $\varphi$  subtended at the driver's retina, and its derivative, the angular rate  $\dot{\varphi}$ , commonly referred to as *looming*<sup>36,42,43</sup>. The model then infers back the kinematic state from those observations, where – with increasing distances – equal variations in  $\varphi$  lead to increasingly larger variations in the inferred distance and therefore increasing uncertainty about the kinematic state of the other agent. In addition, the agent's perception accuracy is limited: it cannot perceive looming with absolute values below a threshold  $\varphi_0$ <sup>44,45</sup>. Due to this threshold, the agent is unable to detect small relative velocities at long distances; our model thus incorporates this looming threshold as one possible mechanism behind delays in recognizing events like the abrupt braking of a lead vehicle.

### Behavior prediction via norm-conditioned particle filter

After perceiving the environment, the ego agent predicts the other vehicle's behavior (Fig. 2b). For this, our model uses a sample-based approach<sup>35,46</sup> based on a particle filter. Specifically, the samples representing the agent's updated belief about the other vehicle are propagated forward by the *generative model* (in our case, using a bicycle model with additive Gaussian noise). In this way, the model generates multiple distinct kinematically plausible future trajectories of the other vehicle, including the ones that could potentially lead to a collision. This enables the model to anticipate rare, long-tail behaviors of the other vehicle and respond to them appropriately during collision avoidance, but yields overly pessimistic predictions in non-conflict situations. For instance, widely dispersed predictions could lead the agent to anticipate any vehicles in the adjacent lane potentially encroaching into its lane, incentivizing unnecessary evasive actions.

Thus, relying solely on kinematic likelihood for the generation of these samples would lead to overly cautious behavior. In reality,



**Fig. 1 | An illustration of the key principles of our active inference model in the opposite-direction lateral incursion scenario.** Upper left panel: The modeled agent (at the bottom, facing forward) is continuously generating beliefs (represented by the arrows) about how the current driving situation will play out and what future observations it will make as a consequence of its own actions and the actions of the other vehicle. In “normal”, non-conflict situations, the situation typically plays out as expected. The agent thus initially believes with high certainty that both itself and the oncoming vehicle will remain in their lanes and pass each other safely. This yields low expected free energy (EFE; green arrows) since the agent’s expected (and thus preferred) observations (e.g., making progress, avoid collisions) are

aligned with the actual observations. Upper right panel: When the oncoming vehicle suddenly encroaches into the driver’s lane, the EFE increases rapidly (as indicated by the red arrows) because this new situation is likely to result in an observed collision in the future, which is strongly disfavored by the agent. In this situation, the agent’s belief about the oncoming vehicle’s future trajectory also becomes more uncertain, because the other vehicle can no longer be trusted to follow established traffic norms. The agent thus needs to imagine a better future with lower EFE -- in this case, a scenario where it avoids collision by swerving left (lower left panel) and takes action to make that better future come true (lower right panel).

human drivers typically must assume that other road users follow traffic rules<sup>47,48</sup> and other social norms for driving<sup>49</sup> to allow for efficient interaction. To reflect this, our model’s sampling process is biased towards trajectories that are not only kinematically feasible but also *normatively likely*, that is, aligned with societal normative expectations on how to behave on the road<sup>50</sup>. We refer to this biased sampling process as a *norm-conditioned particle filter*. In effect, the model’s belief about the other vehicle’s future trajectory is initially constrained by normative expectations, such as that vehicles typically remain within their lane unless there is evidence to suggest otherwise. At the same time, the normative probability is bounded from above: trajectories where norm compliance stays constant or improves over time compared to the starting state are assigned an equal normative likelihood, while those with decreasing norm compliance are deemed less likely. Thus, when another vehicle unexpectedly causes a conflict (for example, by encroaching into the ego vehicle’s lane), the model relaxes the assumption of that vehicle’s future norm-compliance as well, sampling norm-violating trajectories as long as they remain kinematically plausible.

While there is generally a large set of normative expectations dependent on cultural influence, we chose to only implement specific norm-influencing behaviors in each scenario (e.g., staying in one’s lane or yielding to the right of way, see Supplementary Fig. 1). A generalized, scenario-independent implementation of such norms is feasible, but with the goal of computational efficiency was not pursued here.

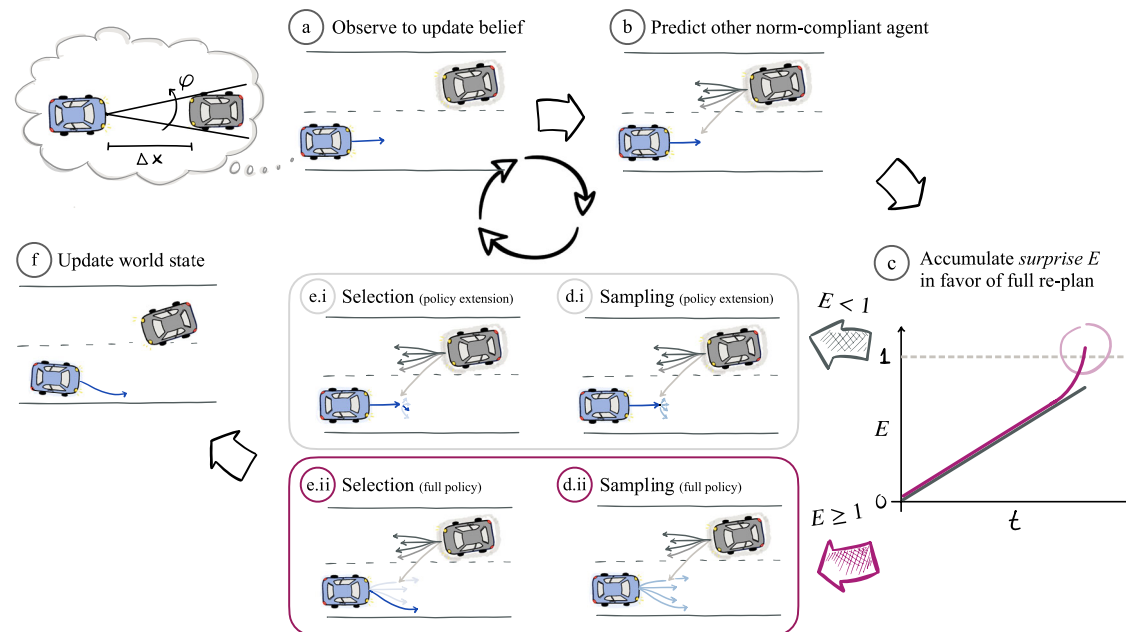
### Re-planning full policy based on accumulating surprise

In the next step, based on the set of previously predicted possible trajectories of the other vehicle, the agent considers its own policy,

evaluating its current suitability by calculating the *residual information*<sup>51</sup> of the *pragmatic value*, which can be seen as a measure of surprise. This surprise signal – the (scaled) negative pragmatic value associated with the currently selected policy – essentially indicates how unsuitable the current policy is given the model’s preferred observations and how the situation is expected to develop. The model assumes that the driver continuously accumulates this surprise signal as evidence in favor of full policy re-plan (Fig. 2c). On each time step, if the accumulated evidence has not reached a predefined threshold, the agent considers it sufficient to continue with the current policy, and extends this policy one time step into the future (Fig. 2d.i and e.i). However, if enough evidence in favor of the current policy being unsuitable is accumulated, a set of completely new policies is sampled and a new full policy is selected (Fig. 2d.ii and e.ii). The first action from either the extended or the new full policy is then applied to update the environment for the next time step (Fig. 2f). This mechanism, where policy re-planning is contingent on the accumulation of evidence, thus provides a way to represent the dynamics of human decision making and response timing in the model.

### Constrained policy sampling

When proposing new candidate policies, either for extension (Fig. 2d.i) or full re-plan (Fig. 2d.ii), the model uses the *cross-entropy method*<sup>52</sup>. With this method, candidate actions are iteratively resampled, increasingly focusing on the most promising areas of the action space. To represent humans’ bounded capacity for planning under time pressure<sup>3,33–36</sup>, we limit the number of evaluated policies in this process. Importantly, the sampled acceleration values are constrained to reflect that humans operate the gas and brake pedals with one foot.



**Fig. 2 | Main functional components of our active inference collision avoidance model.** At each time step, the ego agent observes the world state, updates its belief about the world (a), and makes probabilistic predictions about the actions of the other agent (where future trajectories that adhere to traffic norms are prioritized) (b). At every time step, the agent accumulates evidence (as measured by surprise associated with the current state of the world) on the unsuitability

of the current policy (c). If the accumulated evidence reaches a threshold, the agent re-plans the entire policy, selecting the one that minimizes the expected free energy (EFE) (d.ii and e.ii). Otherwise, it continues using its current policy in an extend form -- i.e., only sampling and then selecting the policy's last time step (d.i and e.i). The first action in the new or extended policy is then applied to the environment (f).

These constraints include limiting applied jerks (as humans cannot press or release pedals instantaneously) and enforcing a constant holding time of 0.2 s at acceleration  $a_0 \leq 0 \text{ ms}^{-2}$  during transitions in both directions between acceleration and deceleration (as humans cannot move their foot between pedals instantaneously).

### Policy selection via expected free energy minimization

Among the sampled candidate policies, the model aims to find the policy that minimizes *expected free energy* (EFE) – the cornerstone of the active inference framework (Fig. 2e.i and e.ii). In our model, the pragmatic value part of the EFE is maximized by (1) maintaining a desired longitudinal velocity, (2) staying within the current lane (avoiding unnecessary or unsafe lane changes) and on the road, (3) minimizing control inputs (avoid harsh braking or steering), (4) preventing collisions or at least reducing their severity (measured by the relative impact velocity), and (5) maintaining sufficient safety margins, avoiding situations where collisions may become unavoidable (such as when following a vehicle too closely without a sufficient stopping distance). Finally, maximizing epistemic value encourages 6) policies that are expected to result in a wide variety of reliable observations, which may help reducing uncertainty<sup>35</sup>. However, in the present collision avoidance setting, behavior is mainly expected to be driven by pragmatic, rather than epistemic, value since there is not much the driver can do to reduce uncertainty, especially given that we here assume that the other vehicle is non-reactive to the ego agent's actions.

### Model evaluation

We evaluated the model against human data in three scenarios.

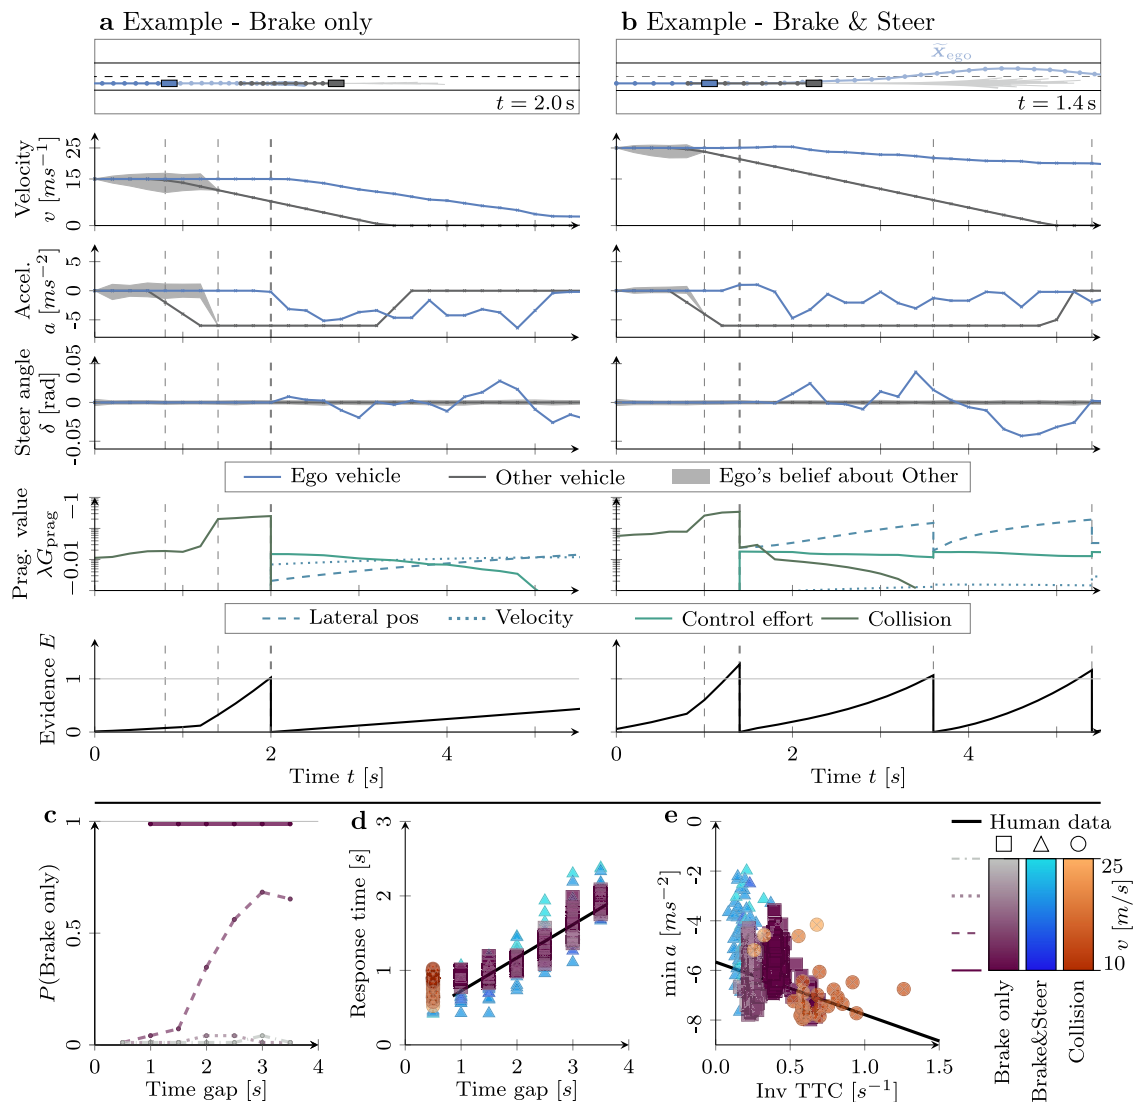
First, in the front-to-rear scenario, we compared the model to the results of a meta-analysis of brake response times<sup>57</sup> in experimental driving simulator and field studies (e.g., Brookhuis et al.<sup>58</sup> and Lee et al.<sup>59</sup>), and an analysis of deceleration magnitudes reported in ref. 14 based on two real-world driving datasets captured across the US (the SHRP2 dataset<sup>60</sup>) and Africa (the ANNEX dataset<sup>61</sup>).

Second, in the opposite-direction lateral incursion scenario, the model was compared to human data from a driving simulator study conducted in the UK<sup>62</sup>, which was conducted in parallel to the present work. Thus, by contrast to the model evaluation for the front-to-rear scenario, which was based on aggregated data, we were here able to run our model on nearly identical scenarios for which the human data was collected, allowing for more detailed comparisons. We manually fitted the model parameters to reproduce human behavior in the first two scenarios; the same parameter values were used for both scenarios.

Third, to test the ability of the model to generalize to new scenarios, we tested it, with the same parameter settings, in an unseen third scenario conducted in Canada<sup>63</sup>, an intersection right-turn-into-path scenario where another vehicle failed to yield the right of way. Similarly to the lateral incursion scenario, we were able to closely reproduce the scenarios faced by the human participants.

**Front-to-rear scenario.** In the front-to-rear scenario, the ego vehicle trails the other vehicle, which is driving in the same lane; both vehicles have the same initial velocity, which was systematically varied in our simulations together with the initial time gap between the vehicles. After a short time of driving at constant speed, the other vehicle starts braking with a high constant deceleration until coming to a stop. Depending on the initial kinematics of the vehicles, the ego agent can avoid a collision either solely by braking (Fig. 3a) or combined braking and swerving (Fig. 3b). The model sometimes avoided a collision by swerving only (i.e., without braking). However, since existing results in the literature on front-to-rear scenarios typically only report on braking response performance, we only compared our model's brake response times to the human response data (i.e., for the analysis of those, steering only responses were not considered).

At lower speeds and larger time gaps, the model typically avoids collisions by braking only. In a representative simulation (Fig. 3a), after the leading agent started braking at  $t = 0.8 \text{ s}$  (first dashed line), it took the model 0.6 s to perceive this (see the ego agent's belief in the



**Fig. 3 | Evaluation of the model in the front-to-rear scenario.** **a** In a scenario where initial pre-conflict kinematics are such that collision can be avoided by braking (initial vehicle velocities  $v_0 = 15 \text{ ms}^{-1}$ , a bumper-to-bumper time gap of 1.5 s), the model typically produces braking-only behavior. The sub-panels visualize the top-down view of the scenario, the time dynamics of both vehicles' velocity  $v$ , acceleration  $a$ , steering angle  $\delta$  and (for the ego vehicle) the components of the pragmatic value  $G_{\text{prag}}$  (scaled with the evidence accumulation gain  $\lambda$ ) that underlie the accumulated evidence towards a full policy re-plan. **b** In a scenario where braking on its own might not be sufficient to avoid collision ( $v_0 = 25 \text{ ms}^{-1}$ , 1 s time gap), the model selects a policy that involves both braking and steering towards the opposite lane (see the model's selected policies in the topmost plots). **c** The likelihood of the ego vehicle deciding to brake and stay behind the leading other

vehicle (without any steering) as a function of time gap, with the four lines representing four different velocities (ranging from  $10 \text{ ms}^{-1}$  to  $25 \text{ ms}^{-1}$ ). This data includes those samples where the agent is only steering. **d** The relationship between the time gap and the brake response time of the ego agent in 896 model simulations (see Supplementary Information 3.1). The regression line based on a meta-analysis of human brake response times<sup>57</sup> is shown for reference. **e** The relationship of the inverse time-to-collision (TTC) at brake onset (defined as  $\frac{\max(0, v_{\text{ego}} - v_{\text{ov}})}{\Delta x}$ ) and the lowest observed deceleration. The regression line based on the human deceleration magnitude data<sup>14</sup> is shown for reference. Video replays of the example simulations are available in the Supplementary Information as Supplementary\_Movie\_1.mp4 (**a**) and Supplementary\_Movie\_2.mp4 (**b**).

acceleration plot). At that point ( $t = 1.4 \text{ s}$ , second dashed line), the model recognized the imminent conflict resulting from this change in perceived behavior of the other vehicle, as can be seen by the sudden decrease (increase in negative value) in the collision part of the pragmatic value. This leads to a corresponding increase in the rate of accumulation of the evidence for a re-plan (i.e., surprise). However, it still took another 0.6 s until the accumulated evidence  $E$  reached the threshold, delaying the first noticeable agent reaction to  $t = 2 \text{ s}$  (third dashed line), where the agent started to execute a new policy – emergency braking. This policy was selected because it allowed the agent to drastically reduce its EFE via reaching future states with high pragmatic value, in particular thanks to low collision probability. This policy was preferred over the alternative policy of braking and steering

into the adjacent lane because the pragmatic value associated with lane changing was more negative than that associated with lower velocity, as the desired (corresponding to the initial) velocity is comparatively low. However, even under this policy, at  $t = 2 \text{ s}$ , the acceleration has still not changed, due to the pedal constraints. Consequently, the first deceleration of the model can be observed at  $t = 2.2 \text{ s}$ , resulting in a total response time of 1.4 s.

Smaller time gaps and higher initial velocities can make it kinematically impossible to avoid a collision by braking only. In such situations, the model typically opts to steer and brake at the same time (Fig. 3b). Here, due to the smaller initial distance, the model's perception delay was noticeably shorter (being only about 0.2 s, with the modeled agent perceiving the braking at  $t = 1 \text{ s}$ , the first dashed line).

However, due to the more challenging initial kinematics, it was more difficult for the model to avoid collision, despite the shorter perception delay. In particular, the model switched between different policies three times. Initially, the model recognized the need to slightly brake and then steer around the other vehicle: already at the time of the first re-plan ( $t = 1.4$  s, second dashed line), the collision component of EFE was substantially reduced by the chosen policy. This policy was mostly sufficient in improving the collision component of the pragmatic value, but as visible in the top panel (planned trajectory  $\bar{X}_{ego}$ ), the agent intends to stay in between lanes for an extended period. This led to two further re-plans ( $t = 3.6$  s and  $t = 5.4$  s, third and fourth dashed lines respectively) to adjust the trajectory. However, even having passed the other vehicle already, the model is not completely successful, with a noticeable lateral position component remaining in the pragmatic value. As the optimal solution of driving in the center of the current lane on a free road should have maximized the pragmatic value, this remaining lateral error is likely the result of the cross-entropy method not identifying the optimal policy. This represents a bounded planning capacity which, as discussed above, is also present in humans<sup>13,53–56</sup> and thus an intended feature of our model.

By systematically varying the initial kinematics of the vehicles, we assessed the model behavior across a spectrum of front-to-rear scenarios. Due to the stochasticity of the model (foremost in behavior prediction and policy sampling), for each of the 28 sets of initial conditions, we ran the simulation of the front-to-rear scenario 32 times. Based on these simulations, we analyzed the model's chosen evasive maneuver, response time, and deceleration magnitude and compared it to human data reported in the driver behavior literature. Figure 3c shows the probability that the model will brake only (without steering) for different velocities and time gaps. Existing studies on how human evasive maneuver choice depends on scenario kinematics in front-to-rear scenarios are limited. Thus, we compared our model results to the few related studies that exist, but these do not provide sufficiently detailed analyses for rigorously evaluating our model's evasive maneuvering decisions. These studies generally suggest that evasive maneuver decisions, similar to response timing, depend on the scenario kinematics (see ref. 62 for a review). Avoiding collisions is generally kinematically more feasible by braking at lower speeds and by steering at higher speeds<sup>64</sup> and human drivers tend to swerve into an adjacent lane if a kinematically viable escape path exists<sup>62,65–67</sup>. In line with these general observations, our model prefers braking at lower velocities and favors combined swerving-and-braking or swerving only responses at higher velocities, regardless of the time gap (Fig. 3c). At medium velocities ( $15 \text{ ms}^{-1}$ ) the model behavior is sensitive to the gap between the vehicles: the longer the time gap, the higher the probability of a brake-only maneuver. This pattern reflects a trade-off between the components of the pragmatic values associated with collisions, control effort (harshness of braking and steering), and lateral position. In more urgent scenarios (short time gaps), the model prefers swerving due to the low pragmatic value of harsh braking. Conversely, in less urgent scenarios (larger time gaps), the model opts for braking since moderate braking avoids collision without incurring the significant reduction of pragmatic value due to leaving the lane (while the deviation from the desired velocity is still tolerable at those medium speeds). These specific dependencies of evasive maneuver choice on scenario kinematics have, to our knowledge, not been empirically tested and represent interesting quantitative predictions from our model.

Regarding brake response times, empirical studies have consistently found a strong dependence on scenario kinematics<sup>2,3,6,14,57,68,69</sup>. Specifically, more kinematically urgent scenarios (e.g., with a small initial time gap and/or hard lead vehicle braking) lead to shorter response times while less urgent scenarios lead to longer response times, with approximately a linear relationship between measures of urgency (e.g., the initial time gap) and mean response time<sup>6,57</sup>. Our

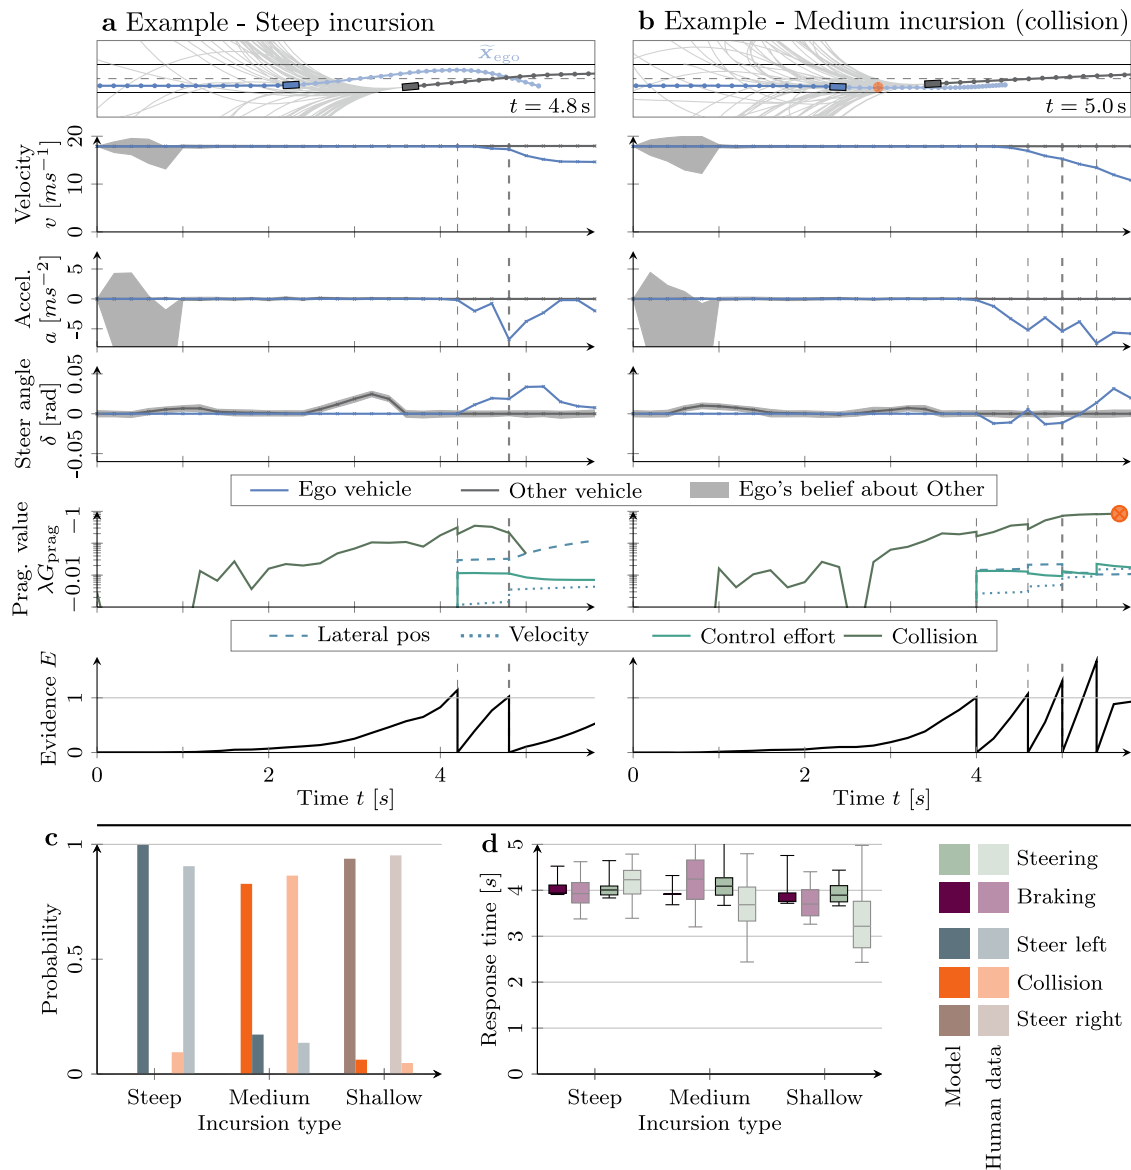
model produced behavior that is remarkably consistent with these findings (Fig. 3d). Interestingly, compared to the evasive maneuver decisions, velocities do not seem to significantly influence brake response times, nor does the final chosen evasive behavior. Furthermore, model response times remain consistent (approx. 1 s) over the range of short time gaps (0.5 s to 1.5 s). This last result could be a consequence of the evidence accumulation mechanism requiring a certain minimum time to trigger a re-plan, even at a very high rate of incoming evidence. Relatedly, the lack of fast enough responses at short time gaps is a likely explanation for the collisions observed in the model (only at the shortest time gap of 0.5 s). However, because the data reported in the literature does not cover such short time gaps, this model prediction remains to be tested in future studies.

While response selection and its timing are comprehensively characterized by the choice of evasive maneuver and response times, response execution is a dynamic process and can be characterized by a variety of metrics. Here we focus on deceleration magnitude, which has been previously shown to depend on the kinematics of the front-to-rear scenario: human drivers brake harder in more kinematically urgent scenarios<sup>14</sup>. Our model captured this phenomenon: the magnitude of its braking increased with the inverse time-to-collision at brake onset (Fig. 3e). This behavior of the model can be explained by the pragmatic value trade-off between avoiding a collision on the one hand and avoiding the effort of hard braking and losing velocity on the other hand. In other words, the model aimed to not brake harder than necessary to avoid collision.

**Opposite-direction lateral incursion scenario.** For evaluating model behavior in the opposite-direction lateral incursion scenario, we used human data from a driving simulator study at the University of Leeds, UK, reported by Johnson et al.<sup>62</sup>. In this study, vehicles initially approached each other in opposite lanes when the computer-controlled vehicle unexpectedly steered toward the participant's lane along a predefined path. The study implemented three kinematic variants differing in incursion "steepness": (1) a *steep* incursion crossing in front of the participant at a sharp angle and allowing a relatively easy escape by steering toward the opposite lane; (2) a *medium* incursion at a shallower angle, heading directly toward the participant; (3) a *shallow* incursion only partially entering the participant's lane, making it possible to find an escape path by steering slightly toward the shoulder (the original study used left-hand traffic, so we are here avoiding directional terms to prevent confusion).

We re-implemented reversed (i.e., right-hand traffic) versions of these three conditions in our setup (see Supplementary Fig. 2), enabling a direct comparison between human and model behavior under nearly identical conditions. Each of those three conditions was repeated 64 times.

An example of our model implemented in the steep incursion condition is shown in Fig. 4a, while Fig. 4b shows an example of the medium incursion condition. When looking at the example simulations, we can observe again that the model needs multiple re-plans before deciding on the final avoidance maneuver. Compared to the front-to-rear scenario, where this is likely caused by the model's bounded planning capacity, in the lateral incursion scenario, this is due to the high uncertainty inherent in the predicted positions of the encroaching vehicle (top panels). In Fig. 4a (steep incursion), a first re-planning is triggered at 4.2 s (first dashed line), but this does not significantly improve the collision component of the pragmatic value. In addition, the model only initiates minor steering maneuvers and small accelerations/decelerations. Together, this suggests that the high uncertainty about the other vehicle's behavior prevents the agent from finding an evasive policy that is better (has lower EFE) than the current policy, instead delaying the decision, thus gaining time to see how the situation develops. With the collision risk unresolved after the first re-plan, surprise quickly accumulates toward a second re-plan at 4.8 s



**Fig. 4 | Evaluation of the model in the opposite-direction lateral incursion scenario.** **a** A steep, unprompted incursion by the other vehicle leads to the model swerving towards the opposite lane. The sub-panels visualize the dynamics of both vehicles' velocity  $v$ , acceleration  $a$ , steering angle  $\delta$  and (for the ego vehicle) the components of the pragmatic value  $G_{\text{prag}}$  (scaled with the evidence accumulation gain  $\lambda$ ) that underlie the accumulated evidence towards a full policy re-plan. **b** In response to a medium incursion by the other vehicle, the model attempts to avoid collision by braking and steering left, but collides at  $t = 5.8$  s (crossed circle marker indicates the location of the collision). **c** The likelihood of the ego vehicle

successfully avoiding the other vehicle by steering left (toward the opposite lane) or right (toward the shoulder), or failing to avoid a collision. This is based on 64 model simulations and 21 to 22 human demonstrations for each of the three incursion types. **d** Brake and steer response times. In panels C and D, human behavior and response times measured in a driving simulator study<sup>62</sup> are shown for reference. The whiskers here show the minimum and maximum observed values. Video replays of the example simulations are available in the Supplementary Information as Supplementary\_Movie\_3.mp4 (a) and Supplementary\_Movie\_4.mp4 (b).

(second dashed line). By this time, the agents are closer, allowing less space for uncertainty to grow and enabling the model to identify a steering policy that mitigates the collision risk (after the second re-plan, the collision component of the pragmatic value no longer significantly contributes to surprise accumulation). The decision to swerve toward the left is a result of both the other vehicle having crossed sufficiently over to the right side to leave a gap for swerving and the model's stronger preference for moving into the opposite lane over leaving the road to the right (similar to Fig. 1). However, it can be seen in Fig. 4a that this trajectory is not optimal. Namely, as can be seen in the top plot, the planned trajectory  $\tilde{x}_{\text{ego}}$  chosen during the second re-plan at  $t = 4.6$  s (second dashed line) over-steers when returning to its own lane and consequently risks leaving the road. Thus, slight

course corrections will likely be performed later (but this would occur outside the current simulation window).

Figure 4b shows an example of the medium incursion scenario where the model is unable to avoid a collision. Again, likely due to uncertainty about the other vehicle's behavior, the first re-plan at  $t = 4.0$  s (first dashed line) mainly results in a sharp deceleration and minor steering towards the right (i.e., towards the shoulder). While steering towards the right might seem suboptimal given the other agent's current path, it aligns with the influence of the norm-conditioned particle filter, which biases the model towards believing that the other agent will likely return to its original lane. At the second re-plan ( $t = 4.6$  s, second dashed line), the model still focuses mostly on continued braking. However, this only marginally improves the

collision component of the pragmatic value, leaving the imminent collision risk unresolved. At the third re-plan at 5 s (third dashed line), the agent keeps braking to reduce impact severity. At this point, the vehicles, which are still on a direct collision course with nearly no lateral offset, got so close that a collision becomes nearly unavoidable. Consequently, even the fourth re-plan ( $t = 5.4$  s, fourth dashed line) fails to find a satisfactory trajectory and thus a collision occurs shortly thereafter at  $t = 5.8$  s.

We compared our model to the human data collected in the driving simulator study by Johnson et al.<sup>62</sup> We closely replicated the scenarios of that study in our simulations (see Supplementary Information 3.2), which allowed direct comparisons between the model and the data with respect to the chosen avoidance behavior and both braking and swerving response times (which were extracted using the same method as in ref. 62; see the “Methods” section). A key finding of Johnson et al.<sup>62</sup> was that the participants’ evasive maneuvering patterns and collision outcomes were strongly determined by the different scenario kinematics in the three incursion scenarios. In the steep scenario, most participants avoided collisions by steering toward the opposite lane, while in the shallow scenario, participants typically steered toward the shoulder, passing the other vehicle on the “inside”. However, in the medium scenario, the majority of participants collided with the oncoming vehicle. Since the urgency (i.e., time to collision at the initial steering of the oncoming vehicle) was constant across scenarios, the high collision rate in the medium case was attributed to greater uncertainty about the other vehicle’s future path, leaving no clear escape route.

As shown in Fig. 4c, our model reproduced these results, reliably avoiding collisions in the steep and shallow scenarios through steering respectively left (toward the center) or right (toward the shoulder). Interestingly, the model also reproduced the human propensity for collision in the medium scenario. Johnson et al.<sup>62</sup> suggested that, conceptually, a key reason why human drivers tended to collide in the medium incursion scenario is the high perceived uncertainty about the oncoming vehicle’s future behavior, which prevents the driver from finding a sufficiently certain escape path. As described above, the current model offers a detailed computational account for the possible mechanisms underlying this phenomenon: the wide spread (uncertainty) in the behavior predictions about the oncoming vehicle, as well as a bias towards expecting that it will return to its own lane due to the norm-conditioning of the beliefs, prevents the model to find an evasive policy in time to avoid collision (see Fig. 4b above). An additional factor behind the observed collision rates was likely also the model’s bounded planning capacity. Increasing the model’s planning capacity – by increasing the number of evaluated policies in the cross entropy method tenfold – resulted in a significant drop in the model’s collision rate (from 82.3% to 51.6% in the medium incursion scenario, and from 6.3% to 3.1% for shallow incursions, for further discussion see Supplementary Information 4).

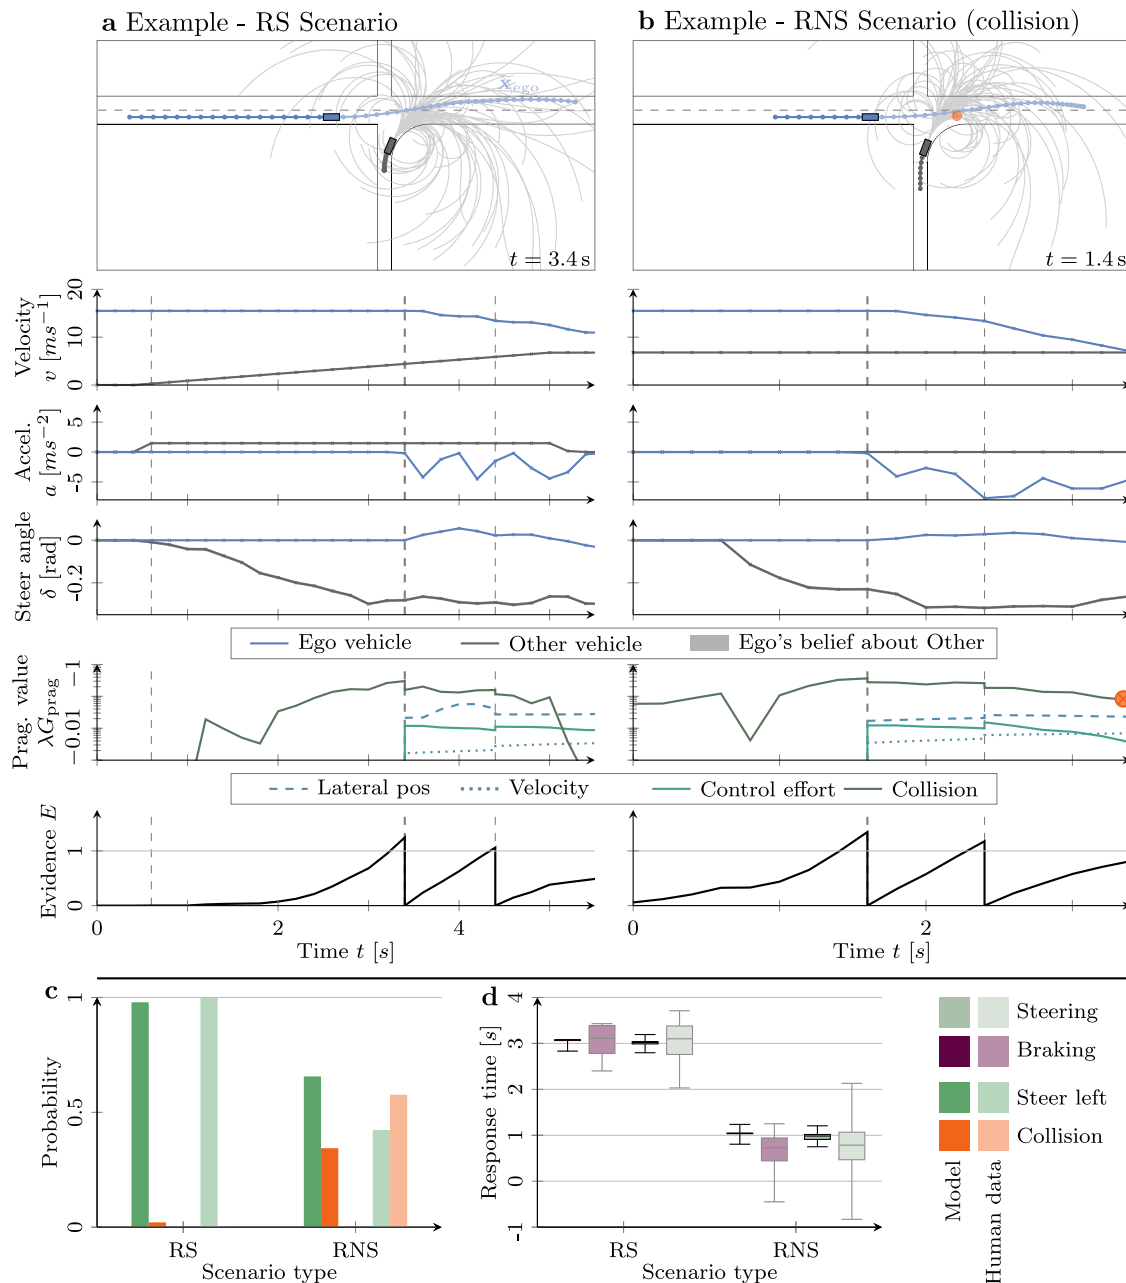
Furthermore, Johnson et al.<sup>62</sup> found that braking and steering response times – measured from initiation of the incursion to the first reaction – were generally long, ranging from 3.5 s to 4.5 s. While braking response times did not vary noticeably with incursion level, steering response did decrease consistently from the steep towards the shallow scenario. Our model generally reproduced the range of response times, with closely matched median values (Fig. 4d), but did not capture the variation in steering response times (especially for the shallow incursion, the model typically steered at least half a second later). One possible explanation may be that many participants avoided collision in the steep scenario primarily by braking before making a final steering adjustment, delaying their recorded steering response. In contrast, the model relied more on steering for collision avoidance, requiring an earlier response. The underlying reasons for these behavioral differences between the model and human drivers are still unclear, but may have to do with varying

individual preferences for braking versus swerving. Further work is needed to see if these individual differences can be reproduced by the model by varying its velocity, lane change, and lateral control effort preferences.

**Intersection scenario.** To test the ability of our model to generalize to more complex scenarios and to unseen data, we evaluated it in an intersection scenario against human data that was not used for tuning model parameters. Specifically, we used human response data obtained in a driving-simulator study conducted in Guelph, Canada, by Ziraldo et al.<sup>63</sup> in the context of collision avoidance at intersections. This study implemented a right-turn-into-path scenario, where the participants approached an intersection with the right of way (green light). A computer-controlled vehicle then approached from the right side, on an intersecting road, and performed a right turn on red into the participant’s lane. Right turns on red are allowed in the US and Canada but the turning vehicle must yield; the participant’s vehicle therefore had the right of way. Two variations of this scenario were implemented: In the right-turn-stopped (RS) scenario variant, the other vehicle initially waited at the stop line in front of the intersection and began to accelerate when the modeled agent’s time-to-arrival was approximately 4 s. In the right-turn-not-stopped (RNS) scenario variant, the other vehicle was moving towards the intersection at a constant speed and was controlled such that, under a constant-velocity assumption, it would cross the stop line when the modeled agent was roughly 1.5 s from entering the intersection. In both scenarios, there was a parallel lane to the left of the modeled agent, allowing it to steer left to avoid a collision. A detailed version of our implementation of these scenario variants can be found in Supplementary Information 3.3. While Ziraldo et al.<sup>63</sup> had each of their 52 participant face three of those scenarios, we only consider the first responses (26 per scenario type), as significant changes in response times and collision rates between repetitions indicated an anticipation of the conflict scenario. In total, we ran each of those scenario variants 96 times, varying the initial velocity of the modeled agent within the limits specified by Ziraldo et al.<sup>63</sup>.

In the RS scenario (Fig. 5a), after the other vehicle begins to accelerate at  $t = 0.6$  s (first dashed line), several time steps pass before the model considers the collision risk to be significant (as reflected in the collision penalty component of the pragmatic value). As evidence (i.e., surprise) thereafter gradually accumulates, the model reaches the re-planning threshold at  $t = 3.4$  s, selecting a policy that combines braking with a leftward steering adjustment. However, as there is still considerable uncertainty about the other vehicle’s future actions – particularly since the other vehicle is already violating normative expectations by not yielding, reducing the bias introduced by norm conditioning – the inferred collision risk decreases only slightly. This necessitates an additional re-planning step one second later (third dashed line). Shortly afterward, the ego vehicle passes the other vehicle, resolving the conflict.

In the RNS scenario (Fig. 5b), the situation becomes critical earlier, as the other vehicle is approaching at constant speed. Since the vehicle does not exhibit any observable braking to yield at the stop line, as would be normatively expected, the model considers it increasingly likely that the other vehicle will encroach into its path. As a result, surprise accumulates based on the increasing collision likelihood. At  $t = 1.6$  s (first dashed line), the ego vehicle initiates braking and applies a slight leftward steering adjustment. However, this response is largely insufficient to reduce the collision risk (the collision component of the pragmatic value does not change substantially). While the model generally predicts the other agent to make a right turn, the high uncertainty in this prediction – given the norm violations by the other vehicle, it is mostly based on kinematic likelihood – leaves no clear safe escape path. The reaction is also too late for braking (although it is attempted) to prevent a collision. Additional re-planning at  $t = 2.4$  s

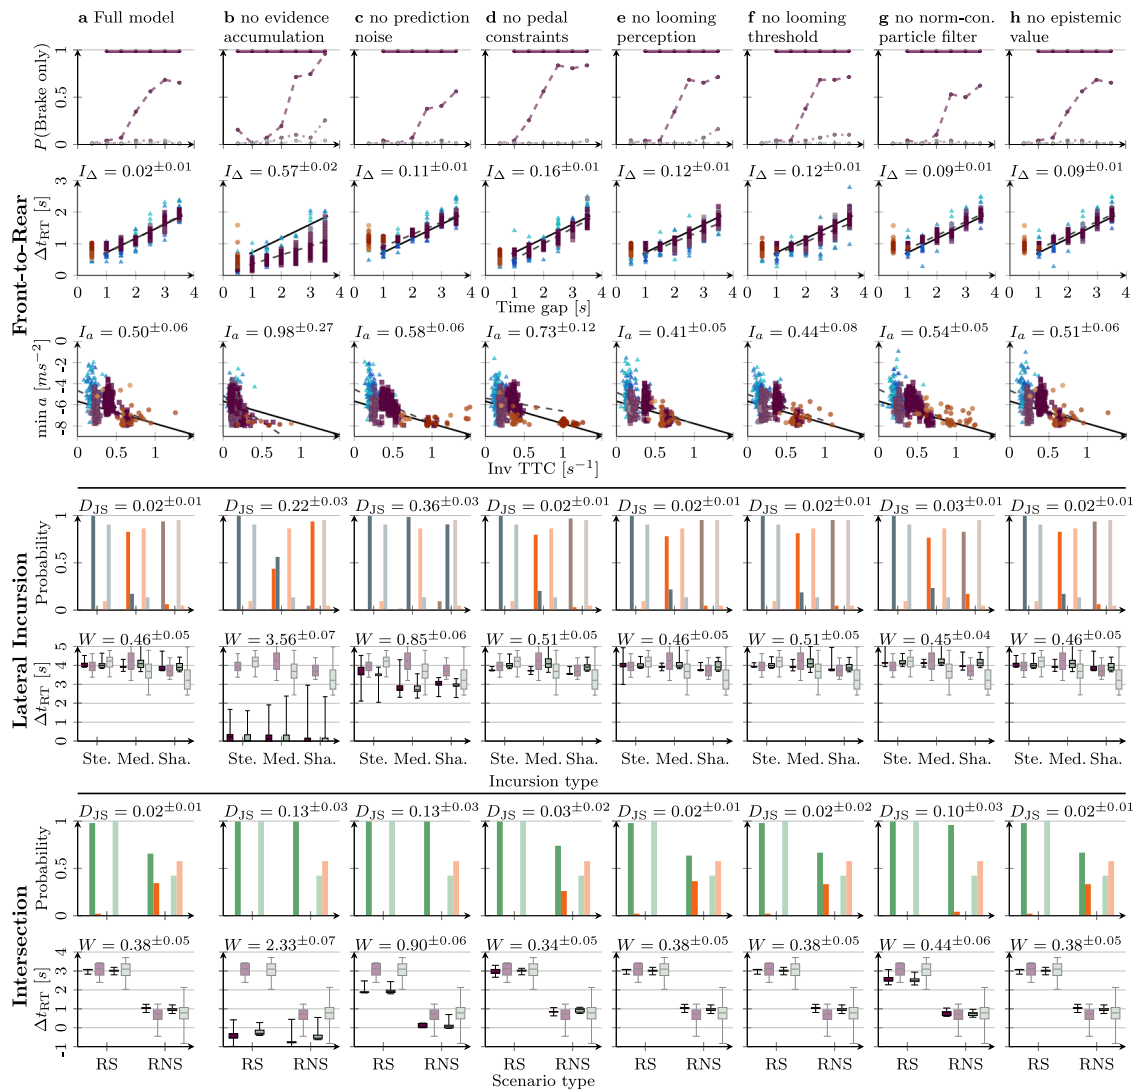


**Fig. 5 | Evaluation of the model in the intersection scenario. a)** Right-turn-stopped (RS): A vehicle initially waiting at the stop line performs a right turn on red without yielding; by accelerating and turning into the ego vehicle’s lane it forces the model to steer left into the parallel lane to avoid a collision. The sub-panels visualize the dynamics of both vehicles’ velocity  $v$ , acceleration  $a$ , steering angle  $\delta$  and (for the ego vehicle) the components of the pragmatic value  $G_{\text{prag}}$  (scaled with the evidence accumulation gain  $\lambda$ ) which constitutes the accumulated evidence towards a full policy re-plan. **b)** Right-turn-not-stopped (RNS): Another agent approaches the intersection with a constant speed, making a right turn into the model’s lane without yielding or even slowing down, leading to a collision at  $t = 3.4$  s

(crossed circle marker indicates the location of the collision). **c)** The likelihood of the ego vehicle successfully avoiding the other vehicle (by steering to the left or braking) or colliding with it, based on 96 model simulations and 26 human demonstrations for each of the two scenario variants. **d)** Brake and steer response times. In panels C and D, human event outcomes and response times obtained in a driving simulator study<sup>63</sup> are shown for reference. The whiskers here show the minimum and maximum observed values. Video replays of the example simulations are available in the Supplementary Information as Supplementary\_Movie\_5.mp4 (a) and Supplementary\_Movie\_6.mp4 (b).

(second dashed line) likewise fails to avert the collision ( $t = 3.4$  s). Had the model opted for a more extreme steering maneuver at the first re-planning point, it could have avoided the collision. However, in this specific simulation, the added pragmatic value of such an extreme maneuver was unclear at that point due to the high uncertainty about the other agent’s future trajectory. Nevertheless, as the policy sampling process is stochastic, the model can occasionally opt for early steep steering and thus avoid collisions in this scenario.

Compared the model to human responses in the same scenario, we found that our model closely reproduced the main observations reported in ref. 63. As in the human data, the model successfully avoided most collisions in the RS scenario, but often collided in the RNS scenario (Fig. 5c). The slightly lower collision rate of the model in the RNS scenario might be explained by the fact that in this scenario approximately ten percent of human participants collided without any noticeable reaction, which might indicate distraction effects that are



**Fig. 6 | Evaluation of model ablations.** Comparing the full model (a) to variants of the model that systematically exclude its key mechanisms (the columns b–h refer to different model ablations). We consider the front-to-rear scenario (top three rows, see Fig. 3 for the legend), the lateral incursion scenario (middle two rows, see Fig. 4 for the legend), and the intersection scenario (bottom two rows, see Fig. 5 for the legend). In the second and third rows, dashed gray lines depict the *maximum a*

*posteriori* linear regression of the model simulation results, with the corresponding mean absolute error *I* between that line and the ground truth ( $I_{\Delta}$  is given in s, while  $I_a$  is given in  $ms^{-2}$ ). In the fourth and sixth row, we used the Jensen-Shannon divergence  $D_{JS}$ , and in the fifth and seventh row the Wasserstein distance  $W$  (in s) to quantify the difference between the model and empirical distributions. For all metrics, expected values  $\pm$  standard deviation are indicated.

not represented in our model. Next to collision rate distributions, our model reproduced the variation in human reaction times in the two scenarios (Fig. 5d): responses in the more critical RNS scenario were markedly faster than in the RS scenario in both human data and model simulations. The model closely matched the timing of human responses in both conditions, with median reaction times generally not being apart more than 0.2 s. This is especially noteworthy, as it was achieved in an unseen scenario without any parameter tuning. This indicates that the model can generalize to new scenarios and can account for human data that was not used for parameter fitting/tuning. While empirically, this generalizability is limited to interactions with a single target agent, given the strong support for the underlying model components in the literature, it can be hypothesized that the model’s generalizability extends further, although this has to be evaluated in future work. Similar to the lateral incursion scenario (Fig. 4d), the variances of the model’s reaction times are smaller compared to the human data. However, this is most likely simply a consequence of using a single model with fixed parameters for all simulations, thereby ignoring potential individual differences between human participants.

**Evaluating individual model mechanisms**

Comparison to the empirical data (Figs. 3, 4, and 5) revealed that our model captured the key aspects of human collision avoidance performance in three different collision avoidance scenarios. To better understand which of the model mechanisms are essential for capturing the empirical observations, we evaluated the performance of seven simpler models, which systematically excluded the key mechanisms one-by-one (Fig. 6).

**Removing evidence accumulation.** Without evidence accumulation, the model fully re-plans its policy at every time step. Consequently, it starts reacting to changes in the other vehicle’s behavior immediately after detecting them. This resulted in a drastic reduction in response times across all scenarios (Fig. 6b). In addition, the model without evidence accumulation predominantly failed to avoid collisions in the shallow incursion scenario, contrasting sharply with human data and the full model. This was mostly caused by the model initiating braking very early. Given the uncertainty about the other vehicle, neither steering left nor steering right allows safe escape paths, so the agent

continues braking, coming to a complete stop. At the point in time where the certainty about the other agent is low enough for safely choosing either left or right, the agent can physically no longer avoid a collision (as it cannot move laterally when stopped). An example simulation of this scenario can be found in the Supplementary Information as *Supplementary\_Movie\_7.mp4*. The fact that collisions happened mostly in the shallow incursion scenarios is likely because in the medium and steep incursion scenarios, premature braking leaves enough space in front of the ego agent for the other vehicle to move across, which is not the case in the shallow incursion. In this scenario, a collision would still occur if the ego agent stayed in its starting lateral position, which makes braking an insufficient avoidance mechanism. Finally, in the intersection scenario, immediate reaction (including both braking and steering) allowed the reduced model to avoid all RNS collisions, in contrast to both the full model and the human data. Based on those observations, the inclusion of evidence accumulation is critical for accounting for empirical data on human response times in all three scenarios, evasive maneuver selection in the shallow lateral incursion scenario, and collision rates in the intersection scenario.

**Removing added noise in behavior prediction.** The model that uses deterministic predictions about the other vehicle produced response times in the front-to-rear scenario that are only slightly longer than for both the full model and the human data (Fig. 6c), especially given shorter time gaps. By contrast, in the opposite-direction lateral incursion and intersection scenarios, removing the noise led to substantially shorter response times. The reason for this apparent discrepancy is that in the front-to-rear scenario, prediction noise results in ego agent's believing that the other vehicle may brake more aggressively, prompting the model to react earlier. In the other scenarios, however, removing the prediction noise decreases the ego agent's uncertainty about the other vehicle, which enables it to find a safe escape path earlier, leading a faster response.

In the lateral incursion scenario specifically, this leads to earlier steering responses, which, given how these particular scenarios unfolded, turned out to be a more effective avoidance strategy, as evidenced by the absence of collisions in this case. However, this strategy was only more effective after the fact (i.e., with hindsight of how the scenario played out). In another counterfactual scenario, where the oncoming vehicle turns back into its own lane early, swerving could instead result in a collision (see ref. 62 for a discussion of this important point). Prediction noise affected not only response times but also the model's choice of evasive maneuver in the lateral incursion scenario. In particular, the model with deterministic predictions predominantly opted to avoid the other vehicle by steering left regardless of the incursion profile. This allowed the model to avoid collisions in all simulations, but is not consistent with the human data.

Similarly, in the intersection scenario, this model ablation successfully avoided all collisions as well. Taken together, these findings underscore that behavior prediction noise is crucial for accurately reproducing human behavior, particularly in the lateral incursion and intersection scenarios. Specifically, uncertainty in the predictions of the other vehicle's actions seems to be a key explanation for why most human drivers ended up in a collision in the medium lateral incursion scenario and the RNS intersection scenario (as discussed above in relation to Figs. 4b, 5b, respectively).

**Removing pedal constraints.** The model not constrained by pedal switching delays resulted in approximately 0.2 s faster brake response times in all three scenarios compared to the full model (Fig. 6d). This matches the delay of 0.2 s implemented to represent the switch of the driver's foot from the gas pedal to the brake pedal. Although the effect appears simple and straightforward, the implemented pedal constraints are essential for capturing the differences between steering and braking response times observed in human behavior. For example,

in the lateral incursion scenario, our model still slightly overestimates steering response times while underestimating braking response times, a discrepancy that could not be corrected by adjusting the surprise threshold for re-planning, as this would shift all response times uniformly.

**Removing looming perception.** We analyzed a version of the model which directly perceived kinematic variables without relying on looming (Fig. 6e) as well as one which perceived looming but did not have a looming perception threshold (Fig. 6f). In the front-to-rear scenario, both these models differed from the full model and the human data in that they produced faster brake response times at longer time gaps. For the model without the looming threshold, this is due to the fact that greater time gaps (and therefore greater distances) require a larger velocity difference to surpass the looming detection threshold, which delays responses, as it needs more time to accumulate. Furthermore, these results being very similar to the model that did not perceive looming at all (Fig. 6e) indicates that in this scenario, the looming threshold was a key factor behind the ability of the model to capture human response times. In addition, while not specifically analyzed here, this effect is expected to be exacerbated in high-speed freeway scenarios where (given the higher speed) the following distances for a given time gap increase, further emphasizing the importance that the model accounts for this phenomenon. In the other two scenarios, removing the looming threshold (or looming perception entirely) did not have any major effects, as the relevant observations in those scenarios are based on lateral movements.

In summary, models without either looming perception or looming threshold resulted in qualitative differences between the model and human data on brake response times in the front-to-rear scenario. In addition to the strong theoretical support, this highlights the importance of both looming-based perception and looming perception threshold for scenarios with largely longitudinal dynamics.

**Removing the norm-conditioned particle filter.** Without the norm-conditioned particle filter, the model behavior in the front-to-rear scenario changes only slightly compared to the full model (Fig. 6g), showing marginally longer response times. In the lateral incursion scenario, however, the effect is more pronounced, with collisions being less likely in the medium incursion (avoided by leftward steering) and more likely in the shallow incursion. Relative to the full model, this ablation lacks the bias against predicting steeper incursions earlier, which incentivizes steering towards the left. This yields a small advantage in the medium incursion scenario, where avoiding towards the other vehicles lane is typically effective, but increases collision risk in the shallow incursion, where both the full model and humans consider rightward avoidance the better option. In addition, the ablated model also shows slightly longer reaction times in the lateral incursion scenario. This again follows from the lack of norm conditioning, as without it, the model often predicts that the other vehicle will leave the road before posing a collision risk, whereas the full model – biased by the normative filter – believes the other vehicle will stay in its lane. Consequently, in this ablation, with the predicted collision likelihood being lower, surprise accumulates more slowly, and the re-planning threshold is reached later, leading to a slower reaction. In the intersection scenario, removing the norm-conditioned particle filter leads to both faster reaction times and a successful resolution of the conflict in nearly every instance. This is logical, as the expectation that the other vehicle would stop to yield resulted in a delayed reaction. Without this bias, potential collisions are taken more seriously and reacted to earlier.

While these differences already highlight the importance of the norm-conditioned particle filter, it has the strongest impact in benign (non-conflict) scenarios. In such scenarios, the model with the simple kinematics-based behavior prediction mechanism is unrealistically

conservative due to having to plan for all kinematically possible futures of the other vehicle – a problem that is avoided by the model with norm-conditioned particle filter (see Supplementary Information 5 for an example).

**Removing the epistemic component of EFE.** Lastly, we investigated a model variant that ignored the epistemic value when selecting policies. This ablation, with the exception of the reaction times in the front-to-rear scenario, did not result in any significant differences in the model behavior, both in the actual outcome and the derived goodness-of-fit metrics. Consequently, it can be argued that the EFE part of our model could potentially be reduced to pragmatic value alone, which would make the model conceptually similar to a belief-space model predictive control<sup>70</sup>. Nevertheless, the epistemic value does play a role in scenarios where, for instance, observability is limited; there, epistemic value has a substantial effect on model behavior (Supplementary Information 6). This complements the previous investigations of the role of epistemic value in scenarios requiring visual behavior<sup>35</sup>. For these reasons, we treat the version of the model with epistemic value as the “full” model here to allow for generalization to similar scenarios in the future.

In addition, in the lateral incursion scenario, we examined the impact of our model’s bounded reasoning capacity. Specifically, as noted above, increasing the number of policies evaluated during replanning led to a substantial reduction in observed collision rates ranging from roughly one-third to one-half across scenarios. This suggests that bounded reasoning is a crucial component for accurately modeling human collision-avoidance failures. However, since not all collisions were avoided, other factors (namely, prediction uncertainty and the expectation of norm-compliant behavior) also play a significant role in driving such outcomes.

Overall, the seven ablated model variants analyzed above show qualitative differences compared to the full model, and none reproduce human behavior across all three scenarios as well as the full model. This provides further evidence for the importance of all the analyzed model mechanisms in regard to capturing human collision avoidance behavior.

## Discussion

Building on previous work<sup>35</sup>, we here presented an active inference-based framework for modeling human collision avoidance behavior based on first principles. To our knowledge, this is the first published computational model that offers a detailed account of human collision avoidance behavior across multiple conflict scenarios. Our model reproduces many key results in the rich literature on collision avoidance in front-to-rear scenarios, in particular the dependence of response times and braking magnitude on scenario kinematics. The model also generated several detailed predictions on how evasive maneuver decisions (braking vs. swerving) vary with scenario kinematics, which can be tested in future experiments. Furthermore, we compared the model to detailed human behavior data in an opposite-direction lateral incursion scenario with varying kinematics<sup>62</sup> and an intersection scenario with a right-turning vehicle failing to yield the right of way<sup>63</sup>, both obtained in driving simulator studies. This analysis demonstrated that the model can reproduce human evasive maneuver decisions, collision outcomes, as well as response times in different variations of these two scenarios. The model also shows strong generalization across scenarios, with the model tuned solely on the front-to-rear and lateral incursion scenario making accurate predictions of human responses in the intersection scenario.

A key feature of our proposed model is its ability to flexibly represent closed-loop collision avoidance through dynamic replanning of policies, including braking, steering, and accelerating actions, in a way that is qualitatively and quantitatively similar to human collision avoidance behavior. In contrast, existing publicly

available collision avoidance models either implement a “one-shot”, often predetermined, open-loop evasive maneuver<sup>1,5,71</sup> or limited aspects of closed-loop control, such as intermittent braking but no steering<sup>11,72</sup>. Some commercially available models – the Stochastic Cognitive Model (SCM) from BMW<sup>73,74</sup> and the driveBOT model from cogniBIT<sup>75,76</sup> – reportedly implement human-like closed-loop behavior in collision avoidance scenarios. However, although the SCM is openly available, the published evaluations provide only a limited basis for meaningful comparison. For instance, Fries et al.<sup>74</sup> evaluate SCM on only a few collision-avoidance cases and compare mainly velocity-over-time profiles, without a systematic analysis of diverse metrics across scenario variations. For driveBOT, the evaluations are more comprehensive<sup>75,76</sup>, but the lack of available information on the detailed working principles of the model precludes a meaningful comparison with our work.

Our model achieves closed-loop behavior by continuously evaluating evasive policies based on their expected free energy, pursuing the policy with the lowest EFE, and allowing for further re-planning if the chosen policy turns out to not yield the preferred (and hence expected) outcomes (e.g., due to limited planning accuracy or unexpected changes in how the scenario unfolds). This, in effect, implements a constraint satisfaction mechanism where different model preferences are traded against each other in finding the policy with the highest overall combined pragmatic value and epistemic value. We observed several examples of this (Fig. 3) where, for example, the model generally prefers to brake at low speeds (since the loss related to the preferred speed is not as high as the cost of changing lane) and swerve at higher speeds (where the cost of braking in terms of speed loss are higher). This yields detailed predictions of human behaviors in different kinematic situations that could be further tested empirically.

Another important aspect of our model is its ability to dynamically account for uncertainty about the future behavior of other agents. This is particularly important in situations like the lateral incursion scenario and the intersection scenario. For instance, in the former, the driver cannot determine a priori if the other vehicle will continue crossing the ego vehicle’s lane or if it will turn back to its original lane, resulting in a wide range of potential future trajectories of the other vehicle. However, the uncertainty may be reduced as the scenario unfolds, opening up new available escape paths, or *escape affordances*<sup>62</sup>, as was the case in the current steep and shallow incursion scenarios. In such cases, an evasive response that is initiated too early may be premature (and non-optimal) and does not properly reflect human behavior. We saw an example of this in Fig. 6c, where in both the lateral incursion and intersection scenario, the model without prediction noise responded much faster than humans and the full model, thus avoiding collision in all cases. However, such overconfident evasive maneuvering behavior could be detrimental in other scenarios, for instance, if in our lateral incursion scenario, the other vehicle would turn back into its original lane. By representing the uncertainty about the other vehicle’s future positions using a vehicle dynamics (bicycle) model and a noisy particle filter, our model is able to handle such situations and successfully reproduce human behavior.

Furthermore, in particular the intersection scenario showed that it is also important that such noisy predictions are biased towards norm-compliant behavior (Fig. 6g). Specifically, without the prior expectation that other vehicle will follow established rules and norms such as yielding the right of way, the model will predict potential collision scenarios earlier, react faster and resolve the conflict safely. However, in many situations, this leads to overreactive non-human like collision avoidance behavior. We also demonstrated that the *norm-conditioned particle filter* is necessary to avoid overly conservative (and therefore non-human like) behavior in non-conflict driving scenarios (Supplementary Information 5).

In the collision avoidance scenarios simulated in the present study, there is little opportunity for the driver to reduce the

uncertainty about the oncoming vehicle's future trajectory through epistemic actions. In other words, there is little the driver can do to obtain further information that could help predict what the other vehicle will do. Thus, as expected, disabling the epistemic value component of the EFE did not significantly change the model behavior. However, as shown in Supplementary Information 6 and previous work<sup>35</sup>, the epistemic value component in our model does drive policy selection in traffic scenarios with epistemic affordances, that is, situations with opportunities for actions yielding information that may resolve uncertainty. Hence, even if the epistemic value did not have any significant effect on policy selection in the present study, it should still be considered a key aspect of human road user behavior. In particular, in the real world, communicative acts such as honking, flashing headlights, or initially moving to the right in the lane could reduce uncertainty about the other vehicle driver's intent. Since other agents were here assumed to be non-reactive, such aspects were not addressed in the current work; this is a topic for future development of the model.

While active inference models can be implemented computationally in many different ways, we used a stochastic receding-horizon control architecture as the basis for the implementation, following<sup>35</sup>. This shares similarities with prior work in robotics, particularly decision-theoretic approaches that explicitly account for uncertainty in actuation and sensing by propagating belief states into the future, often using particle-based representations, and then selecting actions that optimize expected outcomes based on this projected belief evolution<sup>77–79</sup>. However, our work diverges from these engineering models in two crucial aspects. First, since our goal is to represent human driver behavior rather than designing a robot/vehicle controller, the modeling focuses on capturing key aspects of human driving, such as response timing and evasive maneuver decision making, rather than optimizing performance based on engineering criteria. Second, our modeled agent minimizes EFE (as prescribed by the active inference framework), while traditional state-based reward functions often employed in the control literature lack such a theoretically-grounded objective derived from first principles.

Our model incorporates an explicit mechanism for evidence accumulation to account for human response timing, based on existing models<sup>6,43,80–82</sup>. Evidence accumulation models naturally account for the situation-dependency of response timing in traffic situations: the faster a traffic conflict escalates, the faster the human's response to it refs. 6,14,81. Compared to existing evidence accumulation models, our model introduces some key novel aspects. In particular, in contrast to traditional models that accumulate perceptual evidence (e.g., looming<sup>43</sup>), our model rather accumulates surprise. Similar ideas have been explored in existing models<sup>3,6,51,72</sup>, based on the notion that a traffic conflict by definition is an extremely rare (and hence surprising) event which always takes place against a "default" expectation of how the situation would normally play out (Fig. 1). This is also the key idea behind the NIEON (Non-Impaired driver with their Eyes ON the conflict) response time model which conceptualizes the onset of the stimulus that human drivers respond to in a traffic conflict as the onset of surprise (i.e., the violation of the initial default expectation)<sup>6</sup>. The current model generalizes this idea in several ways. First, instead of triggering a predefined action, such as braking when the surprising evidence has reached the threshold, the evidence accumulation here triggers the *re-planning* of a new policy, using the EFE-based policy selection mechanism discussed above. Second, the surprise signal that is being accumulated is not just the difference between a predicted and actual signal but rather the negative pragmatic value of the current policy, that is, how "bad" the currently selected default policy is relative to the agent's preferred observations. This provides a straightforward solution to the known problem of how to determine whether a given surprise signal is relevant to the agent<sup>51</sup>.

Importantly, our model accumulates surprise deterministically without accumulation noise. Under this assumption, the accumulation rate can be expressed in units of the decision boundary (another key parameter traditionally important in the evidence accumulation literature, interpreted as response caution), and hence we did not need to define a separate decision boundary<sup>83</sup>. Nevertheless, adding such accumulation noise could be seen as a future improvement of our model if we want to better match the high variance observed in human reaction times (Figs. 4d, 5d). Similarly, the addition of surprise decay could also be considered.

While we here model evidence accumulation explicitly, it can in principle be seen as an implicit feature of active inference<sup>84</sup>. However, in practice, it is challenging to accurately represent human response dynamics implicitly in a simplified model like ours. To enable straightforward tuning of the model's response performance to human data, we chose the current explicit modeling of evidence accumulation. As demonstrated by our simulation results (Fig. 6b), the inclusion of this mechanism is essential to account for realistic response times in our model.

The current model is based on the general model predictive control architecture originally developed for the model of routine driving described in ref. 35 which, in turn, is based on the standard discrete-time active inference EFE model for planning agents formulated as a partially observable Markov decision process<sup>25</sup>. The original routine driving model in ref. 35 focused primarily on the role of epistemic value in resolving uncertainty when driving around visual occlusions and during visual time sharing (drawing from the rat-in-a-T-maze example in ref. 25). However, that model was not aimed at describing time-critical behaviors. For instance, in collision avoidance scenarios its responses to sudden stimuli would be overly fast, similar to the version of our model without evidence accumulation (Fig. 6b). Extending the original model<sup>35</sup> with not only evidence accumulation, but also other mechanisms such as looming-based perception and norm-conditioned predictions was essential for capturing human collision avoidance behavior (Fig. 6). Otherwise, the model we reported here retains the key features of the original routine driving model needed for dealing with uncertainty resolution through epistemic action, thus offering a powerful and general computational framework for modeling driving behavior, also beyond collision avoidance. As in ref. 35, we employed engineering tools like the cross-entropy method and particle filters with discrete and continuous variables. However, while relying on such established methods, our model does not reduce to a combination of them; instead, these methods serve as means for implementing specific cognitive mechanisms within our active inference framework.

Our framework is grounded not only in active inference and the free energy principle<sup>27</sup>, but also enactivist approaches to cognitive science<sup>85,86</sup>. Active inference and enactivist cognitive science share the foundational notion that behavior in biological agents results from a self-organizing process geared towards sustaining the agent's integrity (existence) over time. This emphasizes a strong continuity between life and mind<sup>86,87</sup> and offers a natural solution to the problem how the certain aspects of the environment become meaningful and relevant to an agent<sup>88</sup>. Such an existential imperative can be described as the organism striving to maintain itself in a sparse attractive set of preferred agent-environment states, where the attractive set depends on the specific organism in question. Thus, to survive over time, a fish needs to remain in water with the right chemical balance and temperature range, obtain food, avoid being eaten by certain predators etc. This implies that these particular states are meaningful and relevant to the fish, which then acts in order to observe these states with a high degree of certainty, thus looking as if it is minimizing free energy over time. By the same token, we endowed our model with a (highly simplified) set of preferred states that are meaningful and relevant to a human car driver (avoid danger of collision, stay on the road, keep up

progress, avoid too harsh braking/steering) and defined the model such that it acts (selects and executes policies) in a way that maximizes the chance of observing these preferred states.

Related to this, our model also offers an operationalization of the classical notion of *affordance*, traditionally broadly defined by Gibson<sup>37</sup> as what the environment “offers the animal, what it provides or furnishes, either for good or ill” (p. 127). Thus, when our model identifies candidate evasive policies associated with low EFE, these can be understood as affordances in the sense of opportunities for actions that are worth pursuing. The active inference framework offers a further distinction between *pragmatic affordances* yielding unsurprising, familiar and preferred observations and *epistemic affordances* yielding information that can resolve uncertainty about pragmatic affordances<sup>35,89,90</sup> (see ref. 62) for a further discussion about this expanded notion of affordances and its relation to the classical affordance concept in ecological psychology).

Thanks to the generality of active inference, models based on it have been developed for virtually every facet of biologically-based cognition and behavior, in organisms ranging from single cells<sup>87</sup>, plants<sup>91</sup>, individual animals and humans<sup>92</sup> to social and cultural phenomena<sup>93</sup>, which provides a great source of concepts and ideas for modeling different aspects of human road user behavior. Thus, importantly, we see active inference not just as a computational approach to model development, but as a general theoretical framework for understanding road user behavior that can guide modeling at a conceptual level and generate hypotheses for experimental studies. At the same time, the great majority of existing active inference models use relatively simple toy scenarios, thereby limiting its applicability to real-world human behavior. By operationalizing active inference in the context of a dynamically rich task, that is, collision avoidance and evaluating it against human data from both simulated and real-world driving, this study pushes the boundaries of active inference applications toward the modeling of complex human naturalistic behaviors.

While, as discussed above, our model represents a significant advancement in human collision avoidance modeling, it has several limitations that can be addressed in future work. First, in this work, we essentially hand-tuned the model parameters to fit the human data from the two collision avoidance scenarios. Whereas this highlights the model’s generalizability and interpretability, especially considering its good performance on the third (intersection) scenario not used for tuning, such an approach is impractical at scale. Future research should explore systematic parameter optimization methods, for example, using the approach suggested by Wei et al.<sup>94</sup>.

Second, the current model assumes that other agents are non-reactive (i.e., in the generative model’s state transition function, other agents’ states are independent from each other and the ego agent). This assumption, therefore, limits the model to scenarios where interactions are only weakly linked and is unlikely to hold in the general case. For example, norms about priority – such as at an unsignalized intersection – are generally interactive, with one agent’s norm compliance depending on the other agents’ behavior. Thus, future work should extend the model to include reactive agents by implementing interactive norms and communicative or epistemic actions that allow agents to probe or signal intentions. Such extensions would strengthen the model’s applicability and allow simulations of multi-agent interactions under active inference. Moreover, the current model (as well as the original routine driving model in ref. 35) only reasons at a shallow “unsophisticated” level about how actions bring about future observations with varying degrees of uncertainty. Further work could explore the possibilities of incorporating *sophisticated inference*<sup>95</sup>, which enables “deeper” levels of recursive reasoning about how future observations would lead to updated beliefs and corresponding new actions. This may be particularly relevant when modeling interactions between road users.

Third, the model currently only accounts for perceptual limitations in terms of visual angle and looming. However, looming only represents the relative motion of objects moving longitudinally towards the observer and, strictly speaking, only applies to objects located along a straight path of travel relative to the gaze direction of the observer (in our case, this means objects located straight ahead of the ego vehicle). Thus, additional perceptual variables are needed to represent perceptual thresholds on the detection of perpendicular motion (e.g., a pedestrian moving into the street ahead) or representing relative motion of objects that are otherwise perpendicular from the observer’s gaze and heading direction. In addition, the perceptual effects of environmental and stimulus-related factors such as weather (e.g., reduced visibility due to fog, heavy rain or snow), lighting conditions (daylight, dusk, night, effects of headlights, etc.), glare, and stimulus conspicuity would ideally need to be accounted for. Another way in which the model’s perceptual fidelity could be improved is by explicitly modeling the human field of view, including the different functional properties of foveal versus peripheral vision, which can be combined with the visual behavior model outlined in the original model<sup>35</sup>.

Beyond perceptual limitations, the use of more realistic perceptual variables such as the optically specified longitudinal time-to-collision ( $\tau$  in<sup>36</sup>), bearing angle (specifying perpendicular collision course<sup>96</sup>), and global optic flow rate or edge rate (specifying self motion relative to the ground<sup>97,98</sup>) could be further explored. This would allow for defining observation and state spaces in the model, which are more aligned with the ecological information available to human drivers.

Fourth, while mechanisms for dealing with partial observability during occlusions and visual time sharing were explored in ref. 35, further work is needed to develop more comprehensive solutions for representing and reasoning about multiple agents hidden behind occluding objects or appearing outside the driver’s field of view (see also Wei et al.<sup>99</sup>); this could help translate the model to more complex and crowded conflict scenarios.

Fifth, the model’s generalizability beyond the presented scenarios from the Anglo-American cultural context has not been evaluated. Three types of assumptions are especially noteworthy: (1) the low-level control assumption that drivers use the same foot for acceleration and braking, (2) the higher-level assumptions about other peoples behavior encoded in normative likelihoods (for example, the different rules regarding right turns on red light in North America and Europe), and (3) the own preferences defining acceptable behavior (such as infringing on other agents’ lanes). As those assumptions are subject to cultural influences<sup>100,101</sup>, they may not transfer seamlessly to other regions of the world. A future model could be made more robust by making such assumptions explicit and learnable (e.g., by modeling them as additional latent (belief) states or context variables dependent on a country).

Finally, while the current policy sampling and kinematic-based particle filter approach were demonstrated to work well in the current scenarios, this could be a limiting factor in scaling the model to a wider range of scenarios, especially when modeling pre-conflict scenarios with more complex road infrastructures and road user interactions. To overcome this limitation, machine-learning-based data-driven models<sup>102–106</sup> could be leveraged to generate possible futures in our generative model’s state transition function, either stand-alone or in combination with the current kinematics-based predictions, thus potentially further enhancing the generalizability of the model. In principle, such data-driven models, if employed in a closed-loop manner, could potentially serve as an alternative to our proposed active inference model in its entirety, given their clear advantages in inference time – no expensive policy selection process is needed – as well as their superior ability to learn from large human datasets in complex geometric environments. These characteristics, after all,

allowed those models to show remarkable and increasing accuracy in the prediction of routine driving behavior<sup>107</sup>. However, research on evaluating such models in collision avoidance scenarios is lacking. To understand their current capabilities, we tested one of the state-of-the-art machine-learned models (ADAPT<sup>22</sup>) in the lateral-incursion scenario (Supplementary Information 7). We found that the data-driven model did not capture human behavior well, likely due to simplistic learning objectives and under-representation of critical scenarios in the training data. Furthermore, the black-box nature of such models restricts them to an exclusively descriptive/predictive role, without providing insight into the underlying cognitive mechanisms driving human behavior.

Overall, our work contributes to the recent body of research<sup>35,94,99,108</sup> with new evidence that active inference can serve as a general framework for computational road user modeling. Along with the emerging real-world applications of active inference in robotics<sup>109</sup> and artificial intelligence<sup>110</sup>, this work can aid in the development of agents used for the simulation-based evaluation of advanced driver assistance systems and automated vehicles<sup>4,111,112</sup>.

## Methods

### Model principles

Our model follows the architecture of the original model initially formulated by Engström et al.<sup>35</sup>, and combines some of the original elements of that model with multiple new mechanisms. In particular, our model retains the representation of probabilistic beliefs using sets of samples and the correspondingly sample-based estimation of the expected free energy (see equation (8)) as well as the cross-entropy method used for policy selection. In addition to these mechanisms, we introduced looming perception, evidence accumulation, the norm-conditioned particle filter, and the constraints on acceleration profiles (see Fig. 6 for the analysis of the impact of these new mechanisms on model behavior). We also modified the Bayesian belief update (equation (2)) and the calculation of the epistemic value (equation (9)).

In our active inference model, the environment is represented by the *generative process*, with a state of the world  $\eta \in \mathcal{E}$  which is partially observable by the agent. The agent can influence the environment with its actions  $\mathbf{a} \in \mathcal{A}$  (according to the state transition probability  $\hat{p}(\eta'|\eta, \mathbf{a})$ ), resulting in observations  $\mathbf{o} \in \mathcal{O}$ , whose dependence on the world state is described in  $\hat{p}(\mathbf{o}|\eta)$  (the  $\hat{p}$  is used to show these function be part of the *generative process*).

The agent has an internal model of the world, the *generative model*, with the corresponding partially observable state  $\mathbf{s} \in \mathcal{S}$ , and the corresponding state transition function  $p(\mathbf{s}'|\mathbf{s}, \mathbf{a})$  and observation probability  $p(\mathbf{o}|\mathbf{s})$ . As the *generative model* is an abstraction of the real world described by the *generative process*,  $\eta$  and  $\mathbf{s}$  do not necessarily have to represent the same state spaces ( $\mathcal{E} \neq \mathcal{S}$ ). The agent is uncertain about the state of the generative model, represented by the probability density function  $q(\mathbf{s})$  (referred to as the agent's belief about the state  $\mathbf{s}$ ). Following<sup>35</sup>, in our model, the agent represents its belief distributions  $q(\mathbf{s})$  in a non-parametric way as a set of  $N$  samples  $\mathbf{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_N\}$  (i.e., a particle filter).

The agent then tries to minimize its free energy, both in the belief  $q(\mathbf{s})$  it forms (*variational free energy*) and the actions  $\mathbf{a}$  it chooses (*expected free energy*). The action and policy selection is based on the existence of a preference prior  $p(\mathbf{o})$  depicting preferred observations.

### Perception

At time  $t$ , the agent perceives the world state  $\eta_t$ , resulting in the observation  $\mathbf{o}_t$ :

$$\mathbf{o}_t \sim \hat{p}(\mathbf{o}'|\eta_t) \quad (1)$$

With this observation, the agent updates its old belief  $q(\mathbf{s}_{t-1})$  to  $q(\mathbf{s}_t)$ . The active inference framework postulates that this update is based on

the minimization of the *variational free energy*<sup>25</sup>, but in our case, we use standard Bayesian updating:

$$q(\mathbf{s}_t) \propto p(\mathbf{o}_t|\mathbf{s}_t) \mathbb{E}_{q(\mathbf{s}_{t-1})} p(\mathbf{s}_t|\mathbf{s}_{t-1}, \mathbf{a}_{t-1}). \quad (2)$$

Here, the particle filter first advances the individual samples in  $\mathbf{S}_{t-1}$  (representing  $q(\mathbf{s}_{t-1})$ ) using the transition function to get the expected states  $\mathbf{S}_{A,t}$  (approximating  $q_A(\mathbf{s}_t) = \exp(\mathbb{E}_{q(\mathbf{s}_{t-1})} \ln p(\mathbf{s}_t|\mathbf{s}_{t-1}, \mathbf{a}_{t-1}))$ ), with  $\mathbf{s}_{A,t,n} \sim p(\mathbf{s}'|\mathbf{s}_{t-1,n}, \mathbf{a}_{t-1})$  being randomly sampled. Then, a kernel density estimate<sup>113,114</sup> based on  $\mathbf{S}_{A,t}$  is used to approximate  $q_A(\mathbf{s}_t)$  as a Gaussian Mixture Model (GMM). Assuming that  $p(\mathbf{o}_t|\mathbf{s}_t)$  is Gaussian as well,  $q(\mathbf{s}_t)$  can also be represented as a GMM; this allows straightforward sampling of  $\mathbf{S}_t$ . Compared to the approach in the original model<sup>35</sup> of weighting samples  $\mathbf{s}_{A,t,n} \in \mathbf{S}_{A,t}$  with  $p(\mathbf{o}_t|\mathbf{s}_{A,t,n})$  and then resampling based thereon, this approach allowed us to represent large shifts in beliefs outside the extreme values of the initial set  $\mathbf{S}_0$ .

The agent's visual perception in our model is based on looming (lementary Information 2.2) – the optical information immediately available to humans about the relative motion of objects straight ahead in the direction of travel<sup>37</sup>. Specifically, the relative position and motion of objects along the forward path of the agent is perceived in terms of the visual angle  $\varphi$  subtended by the object at the retina of the observer, and its derivative, the angular rate  $\dot{\varphi}$  (looming)<sup>36,42</sup>. Due to the nonlinear relationship between the observer's distance to the observed vehicle (in Fig. 2a referred to as  $\Delta x$ ) and  $\varphi$  (with increasing distances, equal variations  $\delta \Delta x$  lead to increasingly smaller  $\delta \varphi$ ), perception errors for  $\varphi$  and  $\dot{\varphi}$  will lead, respectively, to increasing inference errors and therefore uncertainty about the position and speed of the leading other vehicle at increasing distances. We represent this aspect in the model by setting the noise in  $p(\mathbf{o}|\mathbf{s})$  to be constant across  $\varphi$  and its derivatives.

At a certain distance,  $\dot{\varphi}$  is no longer perceptible to the human eye, which implies a minimal threshold for looming detection<sup>44,45</sup>. To account for this, the model incorporates a detection threshold on  $\dot{\varphi}$  (i.e., for small  $|\dot{\varphi}|$ , the means of the observation probabilities  $p(\mathbf{o}|\mathbf{s})$  for velocities and acceleration were set to the values corresponding to  $\dot{\varphi} = 0$ ). This means that changes in the other vehicle's velocity cannot be perceived by a human driver beyond a certain distance, resulting, for example, in a delayed recognition if the lead vehicle suddenly brakes.

### Behavior prediction

After observing the current state of the world, the agent applies the state transition function  $p(\mathbf{s}'|\mathbf{s}, \mathbf{a})$  starting from its current belief  $q(\mathbf{s}_t)$  in order to generate a predicted belief about the future states of the situation  $\tilde{q}_s(\mathbf{s}_\tau|\boldsymbol{\pi}_t, q(\mathbf{s}_t))$  (with  $\tau \in \{t+1, \dots, t+H\}$ , an arbitrary policy  $\boldsymbol{\pi}$  containing the ego agent's sequence of planned actions, and prediction horizon  $H$ ) of all the agents in the scene, as well as the corresponding belief over possible future observations  $\tilde{q}_o(\mathbf{o}_\tau|\boldsymbol{\pi}_t, q(\mathbf{s}_t)) = \mathbb{E}_{\tilde{q}_s(\mathbf{s}_\tau|\boldsymbol{\pi}_t, q(\mathbf{s}_t))} p(\mathbf{o}_\tau|\mathbf{s}_\tau)$ . During this prediction process, the agent adds noise to the control inputs of the other agents to represent the possibility that they might change their behavior.

Practically, the agent advances the belief  $\tilde{q}_s(\mathbf{s}_\tau|\boldsymbol{\pi}_t, q(\mathbf{s}_t))$  (represented by  $\tilde{\mathbf{S}}_\tau = \{\mathbf{s}_{\tau,1}, \dots, \mathbf{s}_{\tau,N}\}$ ) to the next time step by using the state transition function (in our case based on a bicycle model<sup>115</sup>, see Supplementary Information 2.1) on each of the sample states  $\mathbf{s}_{\tau,n} \sim p(\mathbf{s}_\tau|\mathbf{s}_{t-1,n}, \mathbf{a}_{t-1})$  (similarly to forming the expected belief  $q_A(\mathbf{s})$  in the belief update). The belief  $\tilde{q}_o(\mathbf{o}_\tau|\boldsymbol{\pi}_t, q(\mathbf{s}_t))$  is then represented by the set  $\tilde{\mathbf{O}}_\tau = \{\mathbf{o}_{\tau,1}, \dots, \mathbf{o}_{\tau,N}\}$ , where  $\mathbf{o}_{\tau,n} \sim p(\mathbf{o}_\tau|\mathbf{s}_{\tau,n})$ .

This belief propagation scheme represents *unsophisticated inference* in the technical sense that it does not condition the beliefs about future states on corresponding counterfactual observations. While it would theoretically be possible to use a form of *sophisticated inference*<sup>95</sup> that accounts for such conditioning, we here used the

unsophisticated approach for simplicity, similar to the original model<sup>35</sup> and most existing active inference models.

Importantly, our state transition function  $p(\mathbf{s}'|\mathbf{s}, \mathbf{a})$  function is composed out of two parts, with

$$p(\mathbf{s}'|\mathbf{s}, \mathbf{a}) \propto p_n(\mathbf{s}')\check{p}_o(\mathbf{s}'|\mathbf{s}, \mathbf{a}). \quad (3)$$

Here,  $p_o(\mathbf{s}'|\mathbf{s}, \mathbf{a})$  is the standard kinematic likelihood function (as used in ref. 35), while the *projected normative probability*  $\check{p}_n(\mathbf{s}')$  biases the agent to predict future states which currently, and over the near future, are aligned with the norm-compliant behavior of other vehicles. This represents the human tendency to avoid overly pessimistic predictions about the behavior of other traffic participants by instead assuming that they will adhere to road rules and other established traffic norms<sup>47,48</sup>; for example, assuming that oncoming vehicles will stay in their lane or yield the right of way, without evidence to the contrary, is necessary for fluent driving.

In practice, we define  $\check{p}_n$  with

$$\check{p}_n(\mathbf{s}_t) \propto \min\{p_n(\mathbf{s}_t), 2 \frac{p_n(\check{\mathbf{s}}_{t+1})p_n(\check{\mathbf{s}}_{t+H_n})}{p_n(\check{\mathbf{s}}_{t+1}) + p_n(\check{\mathbf{s}}_{t+H_n})}\} \quad (4)$$

applying the *normative probability*  $p_n$  to the current time as well as short- and medium-term predictions  $\check{\mathbf{s}}$  ( $H_n$  is setting the medium-term prediction horizon) to evaluate their norm compliance. This ensures that the model's predictions are not only biased against actions which lead to immediate norm violations, but also against those that put the other vehicle into a position where norm compliance becomes kinematically less likely over time (e.g., initiating a turn while in a straight lane does not immediately lead to the vehicle leaving the road, but might put the vehicle in a position where staying on the road becomes kinematically impossible). Crucially, the use of the current norm-compliance as an upper bound ensures that once norm-violating behavior is observed, the model no longer places trust into norms as a way of constraining predictions about other agents' behavior. In collision avoidance scenarios initiated by the unexpected, norm-violating behavior of other agents, this allows the model to account for all kinematically possible but low-probability "extreme" (long-tail) behaviors.

This approach – which we call *norm-conditioned particle filter* – helps aligning the model's predictions with realistic driving scenarios while accounting for norm violations when necessary. Even though the norms relevant for our studied scenarios (lane following, yielding the right of way) could in general be implemented globally based on map information (for example a detailed graph), for simplicity and computational efficiency we chose to implement in each scenario only those norms that are relevant to that scenario (Supplementary Information 3.1, 3.2 and 3.3; see also Supplementary Information 2.1 for details on sampling from  $p(\mathbf{s}'|\mathbf{s}, \mathbf{a})$ ).

### Policy sampling

Based on the predicted behavior of the other vehicle, the agent selects its policy  $\boldsymbol{\pi}$ . Here, policies are defined as sequences of  $H$  future actions; following the bicycle model, the actions are defined in terms of acceleration  $a_{\text{long}}$  and steering rate  $\omega$ . These actions are selected using the cross-entropy method<sup>35,52</sup>, which iteratively resamples sets of  $M$  candidate policies  $K$  times, where the resampling focuses around the subset of the  $\beta M$  best samples from the previous iteration. Specifically, the agent takes the mean and variance of this promising subset (ignoring potential correlations), and then samples a new set with size  $M$  from the corresponding normal distribution. In contrast to the original model<sup>35</sup> that randomly picked the final sampling output from the last iteration's set, here we specifically choose the best sample.

To better represent human behavior, the model incorporates a pedal constraint to filter out unrealistic changes in acceleration

(Supplementary Information 2.3). As human drivers generally control acceleration by operating the brake and gas pedal with a single foot and cannot move the foot between pedals instantaneously, the model imposes a constant holding time interval of 0.2 s at  $a_0 \leq 0 \text{ ms}^{-2}$  (similar to<sup>116</sup>), such that if the model transitions from a current acceleration  $a_{\text{long}} > a_0$  to a target acceleration  $a_{\text{long},f} < a_0$ , it must first decelerate to  $a_0$ , maintain that value for 0.2 s, and then proceed to  $a_{\text{long},f}$ . The same applies for acceleration changes in the other direction.

### Policy roll-out and evaluation

To evaluate a policy  $\boldsymbol{\pi}$ , the agent first uses the bicycle model to roll out its possible future states based on  $\boldsymbol{\pi}$ . In the general case, the possible future states of the agent and the other vehicle are coupled, hence the agent rolls out its policies jointly with the predicted futures of the other vehicle. For the reason of simplicity, however, our model assumes that in  $p(\mathbf{s}'|\mathbf{s}, \mathbf{a})$  the other vehicle is unresponsive to the actions of the ego vehicle, allowing us to separate behavior prediction and policy roll-outs.

Active inference then rests on the fundamental assumption that an agent prefers policies which minimize its expected free energy (EFE)  $G$  over the prediction horizon of  $H$  time steps; in other words, the best policy is the one with the lowest EFE. This EFE combines the desire of the agent to, on the one hand, observe itself in states that it prefers (maximize pragmatic value  $g_{\text{pragm}}$ ) and, on the other hand, seek information that may reduce uncertainty about the world (maximize epistemic value  $g_{\text{epist}}$ ):

$$G(\boldsymbol{\pi}_t, q(\mathbf{s}_t)) = \sum_{\tau=t+1}^{t+H} -g_{\text{pragm}}(\tilde{q}_o(\mathbf{o}_\tau|\boldsymbol{\pi}_t, q(\mathbf{s}_t))) - g_{\text{epist}}(\tilde{q}_s(\mathbf{s}_\tau|\boldsymbol{\pi}_t, q(\mathbf{s}_t))). \quad (5)$$

The pragmatic values  $g_{\text{pragm}}$  is defined as

$$g_{\text{pragm}}(\tilde{q}_o(\mathbf{o})) = \mathbb{E}_{\tilde{q}_o(\mathbf{o})} \ln p(\mathbf{o}), \quad (6)$$

while we use the Shannon entropy  $\mathcal{H}$  to calculate the epistemic value<sup>25</sup>

$$g_{\text{epist}}(\tilde{q}_s(\mathbf{s})) = \mathbb{E}_{\tilde{q}_o(\mathbf{o})} D_{KL}[\tilde{q}(\mathbf{s}|\mathbf{o})|\tilde{q}_s(\mathbf{s})] = \mathcal{H}(\mathbb{E}_{\tilde{q}_s(\mathbf{s})} p(\mathbf{o}|\mathbf{s})) - \mathbb{E}_{\tilde{q}_s(\mathbf{s})} \mathcal{H}(p(\mathbf{o}|\mathbf{s})). \quad (7)$$

The first term in Equation (7) is known as the *posterior predictive entropy* and represents the extent to which a policy is expected to yield a variety of different observations. The second term, known as *expected ambiguity*, represents the diversity of observations expected for a given state, that is, the extent to which the observations of that state are unreliable (e.g., due to reduced visibility).

In our particle-based representation, we represent  $\tilde{q}_s(\mathbf{s})$  with the set  $\tilde{\mathbf{S}}$  of  $N$  particles and  $\tilde{q}_o(\mathbf{o}) = \mathbb{E}_{\tilde{q}_s(\mathbf{s})} p(\mathbf{o}|\mathbf{s})$  with the set  $\tilde{\mathbf{O}}$ , resulting in the approximations

$$\tilde{g}_{\text{pragm}}(\tilde{\mathbf{O}}) = \frac{1}{N} \sum_{\mathbf{o} \in \tilde{\mathbf{O}}} \ln p(\mathbf{o}) \quad (8)$$

and

$$\tilde{g}_{\text{epist}}(\tilde{\mathbf{S}}, \tilde{\mathbf{O}}) = -\frac{1}{N} \sum_{\mathbf{o} \in \tilde{\mathbf{O}}} \ln \left( \frac{1}{N} \sum_{\mathbf{s} \in \tilde{\mathbf{S}}} p(\mathbf{o}|\mathbf{s}) \right) - \frac{1}{N} \sum_{\mathbf{s} \in \tilde{\mathbf{S}}} \mathcal{H}(p(\mathbf{o}|\mathbf{s})), \quad (9)$$

under the assumption that the entropy  $\mathcal{H}(p(\mathbf{o}|\mathbf{s})) : \mathcal{S} \rightarrow \mathbb{R}$  can be calculated analytically. The original model<sup>35</sup> had calculated  $g_{\text{epist}}$  by using KDE methods based on  $\tilde{\mathbf{O}}$  to approximate  $\tilde{q}_o(\mathbf{o}) = \mathbb{E}_{\tilde{q}_s(\mathbf{s})} p(\mathbf{o}|\mathbf{s})$ ,

compared to us using  $\frac{1}{N} \sum_{\mathbf{s} \in S} \tilde{p}(\mathbf{o}|\mathbf{s})$ . We opted for this approach, as the original estimation would result in an over-approximation of  $\tilde{q}_o(\mathbf{o})$ , given that we evaluate the KDE exclusively at the center of each kernel. Consequently, we use the following approximation  $\tilde{G}$  for the EFE:

$$\tilde{G}(\boldsymbol{\pi}_t, q(\mathbf{s}_t)) = \sum_{\tau=t+1}^{t+H} \tilde{g}_{\text{pragm}}(\tilde{\mathbf{O}}_\tau) - \tilde{g}_{\text{epist}}(\tilde{\mathbf{S}}_\tau, \tilde{\mathbf{O}}_\tau). \tag{10}$$

When evaluating the pragmatic value, we assume that the preference (i.e., the distribution of desired observations)  $p(\mathbf{o})$  accounts for four aspects of the agent’s objective: a) to maintain its desired longitudinal velocity, b) to minimize the magnitude of control inputs, c) to remain on the road and within its current lane (e.g., avoiding lane markers or opposite lanes), d) prevent collisions, and e) avoiding hazardous situations. Each of these aspects is represented by an individual preference function (Supplementary Information 2.4), which are multiplied together:

$$p(\mathbf{o}) = \mathcal{N}(v_{\text{ego}}|\mu_v, \sigma_v) \mathcal{N}(a_{\text{long, ego}}|0, \sigma_a) \mathcal{N}(\omega_{\text{ego}}|0, \sigma_\omega) p_{\text{lat}}(y_{\text{ego}}) p_{\text{coll}}(\mathbf{o}) p_{\text{safe}}(\mathbf{o}). \tag{11}$$

For  $p_{\text{safe}}(\mathbf{o})$ , the model specifically seeks to avoid states where, under the assumption that the other vehicle begins braking suddenly with a deceleration of  $a_{\text{OV, min}}$ , the ego vehicle would fail to avoid a collision despite initiating maximum braking after a response time of 1 s.

**Surprise-based re-planning**

Our model assumes that on every time step, the agent performs planning incrementally, unless it observes a surprising event, in which case it re-plans the full policy. Practically, on every time step, the agent takes the policy chosen in the previous time step, disregards its first action (as it was already executed) and assumes that the remaining actions constitute all but the last actions of the new policy. The above policy selection mechanism is then used to generate the last action of the new policy.

In parallel with incremental policy updates, on every time step, the agent accumulates evidence in favor of a full re-plan, using the accumulation rate  $\lambda$ :

$$E_t = E_{t-1} + \lambda \epsilon_t. \tag{12}$$

Traditionally, in evidence accumulation literature, the accumulation rate is referred to as the drift rate, which is interpreted as the quality of the incoming evidence<sup>83</sup> and can be linked to the decision-maker’s efficiency of processing the perceptual information<sup>117,118</sup>. Here, we follow Engström et al.<sup>6</sup> in accumulating surprise  $\epsilon \geq 0$ , defined as the residual information of the pragmatic value<sup>51</sup>:

$$\epsilon_t = H \max_{\mathbf{o}} \ln p(\mathbf{o}) - \sum_{\tau=t+1}^{t+H} g_{\text{pragm}}(\tilde{q}_o(\mathbf{o}_\tau | \boldsymbol{\pi}_\tau, q(\mathbf{s}_\tau))). \tag{13}$$

Under this definition, the surprise is the difference between the highest pragmatic value possible and the actual pragmatic value of a policy. For instance, if a policy would result in the most preferred observation, the agent would not accumulate any additional evidence. However, if a policy would lead to an undesired (i.e., a priori unlikely) observation, much evidence for a re-plan would be gained.

If the accumulated evidence is below the threshold of 1, the model then follows the extended policy. Otherwise, it generates a completely new policy, applying the above policy selection mechanism to every action in the policy.

Finally, after the policy is determined, its first action  $\mathbf{a}_t$  is used to update the state of the actual world. This is done using the state transition function of the generative process:

$$\boldsymbol{\eta}_{t+1} \sim \tilde{p}(\boldsymbol{\eta}' | \boldsymbol{\eta}_t, \mathbf{a}_t) \tag{14}$$

**Metrics**

**Front-to-rear scenario.** In this scenario, we used the vehicle trajectory and driver input data to extract brake response times and deceleration magnitudes. To that end, we followed Markkula et al.<sup>14</sup>, extracting response times by fitting a piecewise-linear function to the recorded velocity data; the brake response time was then determined as the time instant when the first constant line switches to one with falling velocity. The slope of this line was used as the estimated deceleration magnitude.

When evaluating the goodness of fit of a model to the ground-truth human data, we used the mean absolute error  $l$ . Given a dataset  $D_0 = \{(x, y) | x \in X_0, y \in Y_0\}$  and a linear approximation of the ground truth  $y_{\text{data}}(x) = a_{\text{data}}x + b_{\text{data}}$  over the support interval  $[x_0, x_1]$ , we first discarded all sample pairs from  $D_0$  for which  $x \notin [x_0, x_1]$ , resulting in  $D$ , after which we computed the residuals

$$E = \{y - y_{\text{data}}(x) | (x, y) \in D\}. \tag{15}$$

We assumed the residuals can be approximated by a linear function

$$e(x) = \alpha x + \beta + \epsilon, \epsilon \sim \mathcal{N}(0, \sigma_E), \tag{16}$$

where  $\sigma_E$  is the empirical standard deviation of  $E$ . We placed Gaussian priors

$$p_0(\alpha) = \mathcal{N}\left(0, \frac{\sigma_E}{2\sigma_X}\right), p_0(\beta) = \mathcal{N}\left(0, \frac{\sigma_E}{2}\right), \tag{17}$$

with  $\sigma_X$  denoting the empirical standard deviation of  $X$ . The posterior  $p(\alpha, \beta | D)$  was then obtained using standard Bayesian linear regression<sup>119</sup>. Defining

$$l(\alpha, \beta) = \frac{1}{x_1 - x_0} \int_{x_0}^{x_1} |\alpha x + \beta| dx, \tag{18}$$

we estimated the posterior mean  $\mu_l = \mathbb{E}_{p(\alpha, \beta | D)}[l(\alpha, \beta)]$  and standard deviation  $\sigma_l = \sqrt{\mathbb{V}_{p(\alpha, \beta | D)}[l(\alpha, \beta)]}$  using Monte Carlo sampling from  $p(\alpha, \beta | D)$  with 10000 samples.

When analyzing response times (Fig. 3d and second row of Fig. 6), we set  $x_0 = 0.9$  s and  $x_1 = 3.6$  s. Meanwhile, for the acceleration error (Fig. 3e and third row of Fig. 6), we set  $x_0 = 0$  s<sup>-1</sup> and  $x_1 = 1.0$  s<sup>-1</sup>.

**Opposite-direction lateral incursion scenario.** In this scenario, we extracted brake and steering response times following Johnson et al.<sup>62</sup>, by interpolating along the acceleration  $a_{\text{long}}$  and steering angle  $\delta$  data to find the time at which  $-1$  ms<sup>-2</sup> and 0.0077 rad (i.e., a steering wheel angle of 5°) are exceeded, to extract brake and steering response times respectively.

When evaluating the model fit of our model variants to the human data in this scenario, we used the Jensen-Shannon divergence (a measure of dissimilarity of two categorical distributions) to compare predicted outcomes and the Wasserstein distance (a proper metric) for reaction times. Given the relatively small size of the ground truth dataset, we applied nonparametric bootstrapping<sup>120</sup> with 10000 resamples to estimate the sampling variability of each metric.

**Intersection scenario.** To extract response times, here we followed the same approach as in the lateral incursion scenario. We only

**Table 1 | An overview of the most essential parameters used in the full model, such as in the Cross Entropy Method (CEM) used for policy selection**

Parameter	Description	Source
State transition function $p(\mathbf{s}' \mathbf{s}, \mathbf{a})$		
$\Delta t = 0.2$ s	Time step size	Model of routine driving <sup>35</sup>
4.2m × 1.72 m	Vehicle size	Matching <sup>62</sup>
$w = 3.65$ m	Lane width	Matching <sup>62</sup>
$\sigma_{a,0} = 3$ ms <sup>-2</sup>	Noise applied to other vehicle's acceleration during belief update	Tuned parameter
$\sigma_{\omega,0} = 0.4575$ ms <sup>-2</sup>	Noise applied to other vehicle's steering rate during belief update	Tuned parameter
$\sigma_a = 0.6$ ms <sup>-2</sup>	Noise applied to other vehicle's acceleration during behavior prediction	Tuned parameter
$\sigma_\omega = 0.0915$ ms <sup>-2</sup>	Noise applied to other vehicle's steering rate during behavior prediction	Tuned parameter
$H_n = 20$	Prediction horizon for the <i>projected normative probability</i>	Tuned parameter
Observation probability $p(\mathbf{o} \mathbf{s})$		
$\phi_0 = 0.00215$ s <sup>-1</sup>	Looming threshold	Taken from <sup>45</sup>
Preference function $p(\mathbf{o})$		
$\mu_v = v_0$	Mean of the preference distribution of the agent's velocity $v$	Model of routine driving <sup>35</sup>
$\sigma_v = 0.5$ ms <sup>-1</sup>	Standard deviation of the preference distribution of the agent's velocity $v$	Tuned parameter
$\sigma_a = 0.1$ ms <sup>-2</sup>	Standard deviation of the preference distribution of the agent's acceleration $a$	Tuned parameter
$\sigma_\omega = 0.02$ s <sup>-1</sup>	Standard deviation of the preference distribution of the agent's steering rate $\omega$	Tuned parameter
$g_{LL} = -15000$	Pragmatic value for leaving the road	Tuned parameter
$g_{LC} = -1000$	Pragmatic value for driving on a lane boundary or in opposing lanes	Tuned parameter
$g_C = -10000$	Pragmatic value after collision with relative velocity of 10 ms <sup>-1</sup>	Tuned parameter
Policy selection		
$H = 30$	Prediction horizon	Model of routine driving <sup>35</sup>
$N = 75$	Number of particles	Model of routine driving <sup>35</sup>
$K = 10$	Number of iterations of policy sampling	Model of routine driving <sup>35</sup>
$M = 100$	Number of sampled policies	Tuned parameter
$\beta = 0.1$	$\beta M$ samples with the lowest EFE are used as the base for the next iteration	Model of routine driving <sup>35</sup>
$\mathcal{N}(0, 5$ ms <sup>-2</sup> )	Distribution to sample accelerations $a_i$ in the first CEM iteration	Model of routine driving <sup>35</sup> (mean) and Tuned parameter (std)
$\mathcal{N}(0, 0.1$ s <sup>-1</sup> )	Distribution to sample steering rates $\omega_i$ in the first CEM iteration	Model of routine driving <sup>35</sup> (mean) and Tuned parameter (std)
$a_0 = -0.1$ ms <sup>-2</sup>	Acceleration applied when no pedal is pressed	Logical consideration
$\lambda = 10^{-5.95}$	Evidence accumulation drift rate	Tuned parameter

For tuned parameters, we normally only tested three to four different parameters, with the drift rate  $\lambda$  being the noticeable exception.

compared our model to the first collision instance shown to each participant in the experiment of Ziraldo et al.<sup>63</sup> to ensure that the underlying data represents natural reactions of human drivers. Importantly, the human steering reaction times were extracted based on steering wheel angles. However, as Ziraldo et al.<sup>63</sup> did not report how this relates to the steering angle of the vehicle (which our model works with), there might be some discrepancy here, as we had to rely on the relationship reported by Johnson et al.<sup>62</sup>. In this scenario we used the same bootstrapped metrics as in the lateral incursion scenario (i.e., Jensen-Shannon divergence for collision probabilities and Wasserstein distance for response times).


**Statistic significance.** To determine whether one model variant provided a statistically significant better fit than another, we computed the signal-to-noise ratio of the difference between models, defined as the mean difference in metric values divided by its estimated standard deviation (using sampling from posterior distributions or bootstrapping as described above). A signal-to-noise ratio exceeding 3 was considered indicative of a meaningful difference, following standard statistical heuristics<sup>121</sup>. Here, the difference is computed as the mean metric difference between models, and under the assumption of independence, the standard deviation of the difference is approximated as the square root of the sum of the individual variances.

## Model parameters


Each of the key model mechanisms has a number of parameters associated with it (see Table 1 for an overview of the 26 most essential parameters). For 12 of them, the exact values have been directly adopted from the literature, in particular the original model<sup>35</sup>. Another thirteen model parameters were treated as free parameters and were manually tuned to qualitatively match the empirical observations.

It was not the purpose of this study to obtain the closest possible fit to human data, so no exhaustive parameter optimization was performed. To analyze the impact that changes of model parameters have on its performance, we conducted a sensitivity analysis (Figs. 7 and 8). Specifically, we varied the thirteen tuned parameters from Table 1 one at a time, while keeping the other parameters fixed. For most parameters, this explored deviation by roughly one order of magnitude, if not otherwise constrained. For each of the new parameter sets, we reran all the simulations in the three scenarios described in the Results section, and calculated the goodness-of-fit metrics (Fig. 6). Using the methods described in the previous Metrics section to estimate the uncertainty of those metrics inherent in the limited number of simulations and human demonstrations, we express those in the corresponding box-plots.

The most consequential parameter is the evidence drift rate  $\lambda$  (Equation (12)), where moderate deviations can lead to large changes in reaction times and therefore substantially worse goodness of fit (the



**Fig. 7 | Sensitivity analysis for selected model parameters related to the state transition function and policy selection (Table 1).** The vertical dashed lines indicate the final values used in the full model. We generally changed  $\sigma_{a,0}$  with the same factor. Meanwhile, in the third column, we ensured that  $\frac{\sigma_a}{\sigma_{a,0}} = \frac{\sigma_{a,0}}{\sigma_{a,0}}$ , while keeping  $\sigma_{a,0}$  and  $\sigma_{a,0}$  fixed. In the box plots, the lower and upper whiskers indicate the 5th and 95th percentiles, respectively. For  $\lambda$  in the rightmost column, note the change in y-axis labels for some metrics.



**Fig. 8 | Sensitivity analysis for selected model parameters related to the preference function (Table 1).** The metrics used are the same, the scales of the y-axes are different than in Fig. 7 for illustrative purposes. For  $g_C$  in the rightmost column, note the change in y-axis labels for  $I_a$ . Vertical dashed lines indicate the final values used in the full model. Similar to Fig. 7, we only changed one parameter at a time; the lower and upper whiskers in the box plots indicate the 5th and 95th percentiles, respectively. While the

goodness of fit metrics  $I_a$  and  $W$  in Fig. 7 can degrade on the order of seconds). Model behavior also deteriorates, as seen in the Jensen-Shannon divergences. This is expected, as for example, a sufficient increase in  $\lambda$  to trigger re-planning at (nearly) every timestep would in principle be equivalent to removing the evidence accumulation mechanism completely from the model, which was shown to be disadvantageous for reproducing human behavior in Fig. 6b. However, our chosen parameter value appears to involve a slight trade-off. In the lateral incursion and intersection scenarios, small increases in  $\lambda$  can improve the accuracy of reaction times while worsening outcome prediction, whereas small decreases show the opposite trend. Meanwhile, the other hand-tuned policy selection parameter, the number of simultaneously evaluated policies  $M$ , had even weaker impact (Fig. 7). There, for  $I_a$  and  $D_{JS}$ , increasing or decreasing  $M$  negatively affects the goodness of fit, but there was no evidence that these differences were statistically significant. Meanwhile, for  $I_a$  and  $W$ , lower  $M$  values did yield better fit to the data, but the improvements were marginal.

Parameters of the state transition function did have a major effect on model behavior (Fig. 7). For  $\sigma_{a,0}$  and  $\sigma_{\omega,0}$  (the noise assigned to the other vehicles' control actions in the generative model's state transition function), we generally observed a U-shaped response across all considered metrics, indicating that deviations from the chosen values led to poorer fit to the data, although not all differences were statistically significant. Similarly,  $\sigma_a$  and  $\sigma_\omega$  (the corresponding noise in the policy roll-out during behavior prediction, second column) chosen for the final model generally yielded the best goodness-of-fit. The horizon for calculating the projected normative probability  $H_n$  was less influential: it only impacted model behavior in terms of front-to-rear reaction times ( $I_a$ ) and the prediction of collision outcomes in the intersection scenario ( $D_{JS}$ ).

For the parameters of the preference function  $p(\mathbf{o})$  (Equation (11)), two primary effects can be observed (Fig. 8). First, the choice of  $\sigma_v$  and  $\sigma_a$  is important for producing a reasonable fit for the deceleration magnitude applied in the front-to-rear scenario ( $I_a$  in Fig. 8). Second, the choice of collision cost  $g_C$  is critical, with slight deviations from the selected value leading to substantial decreases in model performance (Fig. 8). Consequently, fine-tuning of this parameter is of high importance. Varying the remaining parameters did not result in statistically significant changes, suggesting that their coarse calibration is sufficient. However, this is somewhat expected, as the initial trajectories of the ego agents often are aligned with the preferences expressed in those parameters, which mostly influence the specific avoidance maneuver chosen, but not the specific reaction times. This can be seen in the specific metric expressing such behavior ( $I_a$ ), where especially  $\sigma_v$  and  $\sigma_a$  have a significant impact.

Especially notable is that – with parameters tuned only on the first two scenarios without any exposure to the intersection scenario – the chosen parameters perform similarly well in the intersection scenario, suggesting good model generalizability. Taken together, this parameter sensitivity analysis indicates that 1) most of the sensitivity of the model is primarily concentrated in  $\lambda$  and  $g_C$ ; 2) most parameters are robust to minor perturbations; and 3) while not necessarily optimal, our final parameter settings achieve a good fit to human data, with no parameter variations leading to substantially better goodness-of-fit across all three scenarios.

## Data availability

The simulation data generated with our active inference model and its ablations in this study have been deposited in the OSF database (<https://osf.io/g54bu>). As mentioned in the corresponding captions, videos of the results shown in Figs. 3, 4, 5 are provided in the Supplementary Materials. Source data underlying the figures is available as a Source Data file. Source data are provided in this paper.

## Code availability

The code used for the data analysis and modeling in this study is available at <https://github.com/tud-hri/Active-Inference-Collision-Avoidance> with <https://doi.org/10.5281/zenodo.20049511>. The code is provided under a non-commercial license (available at <https://github.com/tud-hri/Active-Inference-Collision-Avoidance/blob/main/LICENSE>) that permits use for research, teaching, personal experimentation, and scientific publication. This includes the use of the licensed materials for benchmarking in academic or applied research publications.

## References

- Bärgman, J., Lisovskaja, V., Victor, T., Flanagan, C. & Dozza, M. How does glance behavior influence crash and injury risk? A 'what-if' counterfactual simulation using crashes and near-crashes from SHRP2. *Transp. Res. Part F Traffic Psychol. Behav.* **35**, 152–169 (2015).
- Engström, J., Markkula, G., Xue, Q. & Merat, N. Simulating the effect of cognitive load on braking responses in lead vehicle braking scenarios. *IET Intell. Transp. Syst.* **12**, 427–433 (2018).
- Bianchi Piccinini, G. et al. How do drivers respond to silent automation failures? driving simulator study and comparison of computational driver braking models. *Hum. Factors J. Hum. Factors Ergon. Soc.* **62**, 1212–1229 (2020).
- Montali, N. et al. The Waymo OpenSim agents challenge. *Adv. Neural Inf. Process. Syst.* **36**, 59151–59171 (2023).
- Olleja, P., Markkula, G. & Bärgman, J. Validation of human benchmark models for automated driving system approval: How competent and careful are they really? *Accid. Anal. Prev.* **213**, 107922 (2025).
- Engström, J., Liu, S. Y., Dinparastdjadid, A. & Simoiu, C. Modeling road user response timing in naturalistic traffic conflicts: A surprise-based framework. *Accid. Anal. Prev.* **198**, 107460 (2024).
- Matusz, P. J., Dikker, S., Huth, A. G. & Perrodin, C. Are we ready for real-world neuroscience? *J. Cogn. Neurosci.* **31**, 327–338 (2019).
- Shamay-Tsoory, S. G. & Mendelsohn, A. Real-life neuroscience: an ecological approach to brain and behavior research. *Perspect. Psychol. Sci.* **14**, 841–859 (2019).
- Carvalho, W. & Lampinen, A. Naturalistic computational cognitive science: towards generalizable models and theories that capture the full range of natural behavior. Preprint at <https://doi.org/10.48550/arXiv.2502.20349> (2025).
- Xue, Q., Markkula, G., Yan, X. & Merat, N. Using perceptual cues for brake response to a lead vehicle: Comparing threshold and accumulator models of visual looming. *Accid. Anal. Prev.* **118**, 114–124 (2018).
- Svärd, M., Markkula, G., Bärgman, J. & Victor, T. Computational modeling of driver pre-crash brake response, with and without off-road glances: Parameterization using real-world crashes and near-crashes. *Accid. Anal. Prev.* **163**, 106433 (2021).
- Guo, C., Wang, X., Su, L. & Wang, Y. Safety distance model for longitudinal collision avoidance of logistics vehicles considering slope and road adhesion coefficient. *Proc. Inst. Mech. Eng.* **235**, <https://doi.org/10.1177/0954407020959744> (2021).
- Siebinga, O., Zgonnikov, A. & Abbink, D. A. A model of dyadic merging interactions explains human drivers' behavior from control inputs to decisions. *PNAS Nexus* **3**, 420 (2024).
- Markkula, G., Engström, J., Lodin, J., Bärgman, J. & Victor, T. A farewell to brake reaction times? Kinematics-dependent brake response in naturalistic rear-end emergencies. *Accid. Anal. Prev.* **95**, 209–226 (2016).
- Wei, R., McDonald, A. D., Garcia, A. & Alambeigi, H. Modeling driver responses to automation failures with active inference. *IEEE Trans. Intell. Transp. Syst.* **23**, 18064–18075 (2022).

16. Li, T., Kovaceva, J. & Dozza, M. Modeling collision avoidance maneuvers for micromobility vehicles. *J. Saf. Res.* **87**, 232–243 (2023).
17. Yuan, Y., Weng, X., Ou, Y., & Kitani, K. M. Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting. In *Proceedings of the IEEE/CVF International conference on computer vision* 9813–9823 (2021).
18. Suo, S., Regalado, S., Casas, S. & Urtasun, R. Trafficsim: Learning to simulate realistic multi-agent behaviors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 10400–10409 (2021).
19. Gu, T. et al. Stochastic trajectory prediction via motion indeterminacy diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* 17113–17122 (2022).
20. Igl, M. et al. Symphony: Learning realistic and diverse agents for autonomous driving simulation. In *Proceedings 2022 International Conference on Robotics and Automation (ICRA)* (IEEE, 2022).
21. Mészáros, A., Schumann, J. F., Alonso-Mora, J., Zgonnikov, A., & Kober, J. TrajFlow: Learning Distributions over Trajectories for Human Behavior Prediction. In *2024 IEEE Intelligent Vehicles Symposium (IV)* 184–191 (2024).
22. Aydemir, G., Akan, A. K., & Güney, F. Adapt: Efficient multi-agent trajectory prediction with adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* 8295–8305 (2023).
23. Ivanovic, B., Song, G., Gilitschenski, I. & Pavone, M. trajdata: A unified interface to multiple human trajectory datasets. *Adv. Neural Inf. Process. Syst.* **36**, 27582–27593 (2023).
24. Schumann, J. F., Kober, J. & Zgonnikov, A. Benchmarking behavior prediction models in gap acceptance scenarios. *IEEE Trans. Intell. Veh.* **8**, 2580–2591 (2023).
25. Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P. & Pezzulo, G. Active inference: a process theory. *Neural Comput.* **29**, 1–49 (2017).
26. Da Costa, L., Parr, T., Sajid, N., Veselic, S., Neacsu, V. & Friston, K. Active inference on discrete state-spaces: a synthesis. *J. Math. Psychol.* **99**, 102447 (2020).
27. Parr, T., Pezzulo, G. & Friston, K.J. *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior* (MIT Press, 2022).
28. Smith, R., Badcock, P. & Friston, K. J. Recent advances in the application of predictive coding and active inference models within clinical neuroscience. *Psychiatry Clin. Neurosci.* **75**, 3–13 (2021).
29. Vasil, J., Badcock, P. B., Constant, A., Friston, K. & Ramstead, M. J. D. A world unto itself: human communication as active inference. *Front. Psychol.* **11**, 417 (2020).
30. Botteman, H. & Friston, K. J. An active inference account of protective behaviours during the COVID-19 pandemic. *Cogn. Affect. Behav. Neurosci.* **21**, 1117–1129 (2021).
31. Harris, D. J. et al. An active inference account of skilled anticipation in sport: using computational models to formalise theory and generate new hypotheses. *Sports Med.* **52**, 2023–2038 (2022).
32. Novický, F., Meera, A.A., Zeldenrust, F. & Lanillos, P. Precision not prediction: Body-ownership illusion as a consequence of online precision adaptation under Bayesian inference. Preprint at <https://doi.org/10.1101/2024.09.04.611162> (2024).
33. Engström, J. et al. Great expectations: a predictive processing account of automobile driving. *Theor. Issues Ergon. Sci.* **19**, 156–194 (2018).
34. Wei, R. et al. *World Model Learning from Demonstrations with Active Inference: Application to Driving Behavior*, in *Active Inference* (Springer Nature Switzerland, Cham, 2023).
35. Engström, J. et al. Resolving uncertainty on the fly: modeling adaptive driving behavior as active inference. *Front. Neurobot.* **18**, 1341750 (2024).
36. Lee, D. N. A theory of visual control of braking based on information about time-to-collision. *Perception* **5**, 437–459 (1976).
37. Gibson, J.J. *The Ecological Approach to Visual Perception: Classic Edition* (Psychology Press, New York, 2014).
38. Luce, R.D. *Response Times: Their Role in Inferring Elementary Mental Organization* (Oxford University Press on Demand, 1986).
39. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
40. Ratcliff, R. & Van Dongen, H. P. Diffusion model for one-choice reaction-time tasks and the cognitive effects of sleep deprivation. *Proc. Natl. Acad. Sci. USA* **108**, 11285–11290 (2011).
41. Markkula, G. Modeling driver control behavior in both routine and near-accident driving. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (SAGE Publications, Sage CA: Los Angeles, CA, 2014).
42. Pekkanen, J. et al. A computational model for driver’s cognitive state, visual perception and intermittent attention in a distracted car following task. *R. Soc. Open Sci.* **5**, 180194 (2018).
43. Tian, K. et al. Explaining unsafe pedestrian road crossing behaviours using a psychophysics-based gap acceptance model. *Saf. Sci.* **154**, 105837 (2022).
44. Hoffmann, E. R. Estimation of time to vehicle arrival-effects of age on use of available visual information. *Perception* **23**, 947–955 (1994).
45. Lamble, D., Laakso, M. & Summala, H. Detection thresholds in car following situations and peripheral vision: implications for positioning of visually demanding in-car displays. *Ergonomics* **42**, 807–815 (1999).
46. Murphy, K.P. *Machine Learning: A Probabilistic Perspective* (MIT press, 2012).
47. Laurent, H., Sangnier, M. & Treibich, C. Traffic safety and norms of compliance with rules: An exploratory study. *Econ. Bull.* **41**, 2464–2483 (2021).
48. Abbas, S., Fatima, S., Sharif, A. & Adnan, S. M. Drivers’ knowledge, attitude and practices towards traffic rules and regulations. *J. Road. Saf.* **35**, 24–31 (2024).
49. Tennant, C. et al. Code, culture, and concrete: Self-driving vehicles and the rules of the road. *Front. Sustain. Cities* **3**, 710478 (2021).
50. Bicchieri, C. *The Grammar of Society*(2006).
51. Dinparastdjadid, A., Supeene, I. & Engstrom, J. Measuring surprise in the wild. Preprint at <https://doi.org/10.48550/arXiv.2305.07733> (2023).
52. de Boer, P. T., Kroese, D. P., Mannor, S. & Rubinstein, R. Y. A tutorial on the cross-entropy method. *Ann. Oper. Res.* **134**, 19–67 (2005).
53. Simon, H.A. et al. A behavioral model of rational choice. *Q. J. Econ.* **69**, 99–118 (1955).
54. Summala, H. in *Modelling Driver Behaviour in Automotive Environments: Critical Issues in Driver Interactions with Intelligent Transport Systems*, ed. by Cacciabue, P.C. (Springer, London, 2007).
55. Oh, H. et al. Satisficing in split-second decision making is characterized by strategic cue discounting. *J. Exp. Psychol. Learn. Mem. Cogn.* **42**, 1937–1956 (2016).
56. Callaway, F. et al. Rational use of cognitive resources in human planning. *Nat. Hum. Behav.* **6**, 1112–1125 (2022).
57. Engstroem, J. et al. Scenario criticality determines the effect of working memory load on brake response time. In *Proceedings of the European Conference on Human Centred Design for Intelligent Transport Systems* (2010).
58. Brookhuis, K. A., de Vries, G. & de Waard, D. The effects of mobile telephoning on driving performance. *Accid. Anal. Prev.* **23**, 309–316 (1991).

59. Lee, J. D., Caven, B., Haake, S. & Brown, T. L. Speech-based interaction with in-vehicle computers: the effect of speech-based e-mail on drivers' attention to the roadway. *Hum. Factors* **43**, 631–640 (2001).
60. Victor, T. et al. Analysis of naturalistic driving study data: Safer glances, driver inattention, and crash risk. <https://nap.nationalacademies.org/catalog/22297> (2015).
61. Engström, J., Werneke, J., Bärgrman, J., Nguyen, N. & Cook, B. Analysis of the role of inattention in road crashes based on naturalistic on-board safety monitoring data. In *Proceedings 3rd International Conference on Driver Distraction and Inattention* (2013).
62. Johnson, L., Engström, J., Srinivasan, A., Öztürk, I. & Markkula, G. Looking for an out: Affordances, uncertainty and collision avoidance behavior of human drivers. Preprint at <https://doi.org/10.48550/arXiv.2505.14842> (2025).
63. Ziraldo, E., Attalla, S., Kods, S. & Oliver, M. Driver response to right turning path intrusions at signalized intersections. *SAE Int. J. Adv. Curr. Pract. Mobil.* **2**, 1623–1632 (2020).
64. Brännström, M., Coelingh, E. & Sjöberg, J. Decision-making on when to brake and when to steer to avoid a collision. *Int. J. Veh. Saf.* **7**, 87–106 (2014).
65. Venkatraman, V., Lee, J. D. & Schwarz, C. W. Steer or brake?: modeling drivers' collision-avoidance behavior by using perceptual cues. *Transp. Res. Rec.* **2602**, 97–103 (2016).
66. Hu, M. et al. Decision tree-based maneuver prediction for driver rear-end risk-avoidance behaviors in cut-in scenarios. *J. Adv. Transp.* **2017**, e7170358 (2017).
67. Sarkar, A. et al. Steering or braking avoidance response in SHRP2 rear-end crashes and near-crashes: A decision tree approach. *Accid. Anal. Prev.* **154**, 106055 (2021).
68. Aust, M. L., Engström, J. & Viström, M. Effects of forward collision warning and repeated event exposure on emergency braking. *Transp. Res. Part F Traffic Psychol. Behav.* **18**, 34–46 (2013).
69. McDonald, A. D. et al. Toward computational simulations of behavior during automated driving takeovers: a review of the Empirical and Modeling Literatures. *Hum. Factors* **61**, 642–688 (2019).
70. Van Den Berg, J., Patil, S. & Alterovitz, R. Motion planning under uncertainty using iterative local optimization in belief space. *Int. J. Robot. Res.* **31**, 1263–1278 (2012).
71. Kusano, K. D. & Gabler, H. C. Safety benefits of forward collision warning, brake assist, and autonomous braking systems in rear-end collisions. *IEEE Trans. Intell. Transp. Syst.* **13**, 1546–1555 (2012).
72. Svärd, M., Markkula, G., Engström, J., Granum, F. & Bärgrman, J. A quantitative driver model of pre-crash brake onset and control. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* (SAGE Publications, 2017).
73. Fries, A., Fahrenkrog, F., Donauer, K., Mai, M. & Raisch, F. Driver behavior model for the safety assessment of automated driving. In *2022 IEEE Intelligent Vehicles Symposium (IV)* (IEEE, 2022).
74. Fries, A., Lemberg, L., Fahrenkrog, F., Mai, M. & Das, A. Modeling driver behavior in critical traffic scenarios for the safety assessment of automated driving. *Traffic Inj. Prev.* **24**, S105–S110 (2023).
75. de Oliveira, L.F.A., Schories, L., Brostek, L. & Meywerk, M. Simulation-based evaluation of a generic autonomous emergency braking system using a cognitive pedestrian behavior model. In *Proceedings 27th International Technical Conference on the Enhanced Safety of Vehicles (ESV)* (2023).
76. Rössert, C., Drever, J. & Brostek, L. Cognitive behavior model replicates road user response timing in naturalistic rear-end traffic conflicts <https://doi.org/10.31219/osf.io/su5kt> (2024).
77. Hubmann, C., Becker, M., Althoff, D., Lenz, D. & Stiller, C. Decision making for autonomous driving considering interaction and uncertain prediction of surrounding vehicles. In *Proceedings 2017 IEEE Intelligent Vehicles Symposium (IV)* (IEEE, 2017).
78. Brechtel, S., Gindele, T. & Dillmann, R. Probabilistic decision-making under uncertainty for autonomous driving using continuous POMDPs. In *Proceedings 17th International IEEE Conference on Intelligent Transportation Systems (ITSC)* (2014).
79. Bouton, M., Cosgun, A. & Kochenderfer, M.J. Belief state planning for autonomously navigating urban intersections. In *Proceedings 2017 IEEE Intelligent Vehicles Symposium (IV)* (IEEE, 2017).
80. Pekkanen, J. et al. Variable-drift diffusion models of pedestrian road-crossing decisions. *Comput. Brain Behav.* **5**, 60–80 (2022).
81. Zgonnikov, A., Abbink, D. & Markkula, G. Should I stay or should I go? Cognitive modeling of left-turn gap acceptance decisions in human drivers. *Hum. Fact.* **66**, 1399–1413 (2022).
82. Schumann, J.F., Srinivasan, A.R., Kober, J., Markkula, G. & Zgonnikov, A. Using models based on cognitive theory to predict human behavior in traffic: a case study. In *Proceedings 2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*(2023).
83. Ratcliff, R., Smith, P. L., Brown, S. D. & McKoon, G. Diffusion decision model: Current issues and history. *Trends Cogn. Sci.* **20**, 260–281 (2016).
84. FitzGerald, T. H. B., Schwartenbeck, P., Moutoussis, M., Dolan, R. J. & Friston, K. Active inference, evidence accumulation, and the urn task. *Neural Comput.* **27**, 306–328 (2015).
85. Varela, F.J., Thompson, E., Rosch, E. *The Embodied Mind, revised edition: Cognitive Science and Human Experience* (MIT Press, 2017).
86. Thompson, E. *Mind in Life: Biology, Phenomenology, and the Sciences of Mind* (Harvard University Press, 2010).
87. Friston, K. Life as we know it. *J. R. Soc. Interface* **10**, 20130475 (2013).
88. Kiverstein, J., Kirchhoff, M. D. & Froese, T. The problem of meaning: the free energy principle and artificial agency. *Front. Neurobot.* **16**, 844773 (2022).
89. Friston, K. et al. in *Affordances in Everyday Life: A Multidisciplinary Collection of Essays* (Springer International Publishing, Cham, 2022).
90. Ramstead, M.J.D. et al. in *Affordances in Everyday Life: A Multidisciplinary Collection of Essays* (Springer International Publishing, Cham, 2022).
91. Calvo, P. & Friston, K. Predicting green: really radical (plant) predictive processing. *J. R. Soc. Interface* **14**, 20170096 (2017).
92. Van de Maele, T., Dhoedt, B., Verbelen, T. & Pezzulo, G. A hierarchical active inference model of spatial alternation tasks and the hippocampal-prefrontal circuit. *Nat. Commun.* **15**, 9892 (2024).
93. Kastel, N., Hesp, C., Ridderinkhof, K. R. & Friston, K. J. Small steps for mankind: Modeling the emergence of cumulative culture from joint active inference communication. *Front. Neurobot.* **16**, 944986 (2023).
94. Wei, R. et al. Learning an active inference model of driver perception and control: Application to vehicle car-following. In *Proceedings IEEE Transactions on Intelligent Transportation Systems* (2025).
95. Friston, K., Da Costa, L., Hafner, D., Hesp, C. & Parr, T. Sophisticated inference. *Neural Comput.* **33**, 713–763 (2021).
96. Chohan, A., Savelsbergh, G. J., Van Kampen, P., Wind, M. & Verheul, M. H. Postural adjustments and bearing angle use in interceptive actions. *Exp. Brain Res.* **171**, 47–55 (2006).
97. Larish, J. F. & Flach, J. M. Sources of optical information useful for perception of speed of rectilinear self-motion. *J. Exp. Psychol. Hum. Percept. Perform.* **16**, 295 (1990).
98. Fajen, B. R. Calibration, information, and control strategies for braking to avoid a collision. *J. Exp. Psychol. Hum. Percept. Perform.* **31**, 480 (2005).

99. Wei, R., Lee, J., Wakayama, S., Tschantz, A., Heins, C., Buckley, C., ... & Rajagopal, R. Navigation under uncertainty: Trajectory prediction and occlusion reasoning with switching dynamical systems. In *2025 IEEE International Automated Vehicle Validation Conference (IAVVC)* 1–6 (2025).
100. Pelé, M. et al. Cultural influence of social information use in pedestrian road-crossing behaviours. *R. Soc. Open Sci.* **4**, 160739 (2017).
101. Özkan, T., Lajunen, T., Chliaoutakis, J. E., Parker, D. & Summala, H. Cross-cultural differences in driving behaviours: A comparison of six countries. *Transp. Res. Part F Traffic Psychol. Behav.* **9**, 227–242 (2006).
102. Salzmann, T., Ivanovic, B., Chakravarty, P. & Pavone, M. Trajectron ++: Dynamically-feasible trajectory forecasting with heterogeneous data. In *European conference on computer vision* 683–700 (2020).
103. Mangalam, K., An, Y., Girase, H. & Malik, J. From Goals, Waypoints & Paths to long term human trajectory forecasting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021).
104. Girgis, R. et al. Latent variable sequential set transformers for joint multi-agent motion prediction. In *Proceedings 10th International Conference on Learning Representations, ICLR 2022* (2022).
105. Shi, S., Jiang, L., Dai, D. & Schiele, B. Motion transformer with global intention localization and local movement refinement. *Adv. Neural Inf. Process. Syst.* **35**, 6531–6543 (2022).
106. Nayakanti, N. et al. Wayformer: motion forecasting via simple & efficient attention networks. In *Proceedings 2023 IEEE International Conference on Robotics and Automation (ICRA)* (2023).
107. Schumann, J.F., Mészáros, A., Kober, J. & Zgonnikov, A. Step: Structured training and evaluation platform for benchmarking trajectory prediction models. Preprint at <https://doi.org/10.48550/arXiv.2509.14801> (2025).
108. Wei, R., McDonald, A.D., Mehta, R.K. & Garcia, A. Active inference models of AV takeovers: relating model parameters to trust, situation awareness, and fatigue. *Hum. Factors* **67**, <https://doi.org/10.1177/00187208241295932> (2024).
109. Lanillos, P. et al. Active inference in robotics and artificial agents: survey and challenges. Preprint at <https://doi.org/10.48550/arXiv.2112.01871> (2021).
110. Mazzaglia, P., Verbelen, T., Çatal, O. & Dhoedt, B. The free energy principle for perception and action: a deep learning perspective. *Entropy* **24**, 301 (2022).
111. Fremont, D.J. et al. Formal scenario-based testing of autonomous vehicles: From simulation to the real world. In *Proceedings 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)* (IEEE, 2020).
112. Uzzaman, A., Islam, M. & Hossain, M. S. Testing autonomous vehicles in virtual environments: A review of simulation tools and techniques. *Control Syst. Optim. Lett.* **3**, 151–157 (2025).
113. Deisenroth, M.P., Faisal, A.A., Ong, C.S. *Mathematics for Machine Learning* (Cambridge University Press, 2020).
114. Fischer, J. & Tas, O.S. Information particle filter tree: an online algorithm for POMDPs with belief-based rewards on continuous domains. In *Proceedings of the 37th International Conference on Machine Learning (PMLR)*, 2020).
115. Polack, P., Altché, F., d'Andréa-Novel, B. & de La Fortelle, A. The kinematic bicycle model: A consistent model for planning feasible trajectories for autonomous vehicles? In *Proceedings 2017 IEEE Intelligent Vehicles Symposium (IV)* (2017).
116. Weber, H., Eckstein, L., Beringhoff, F. & Josten, J. Towards modeling driver performance within crash-relevant scenarios as a virtual reference for the Safety of Automated Vehicles. In *Proceedings 27th International Technical Conference on the Enhanced Safety of Vehicles (ESV)* (2023).
117. Metin, B. et al. Adhd performance reflects inefficient but not impulsive information processing: a diffusion model analysis. *Neuropsychology* **27**, 193 (2013).
118. Reinhartz, A., Strobach, T., Jacobsen, T. & von Bastian, C. C. Mechanisms of training-related change in processing speed: a drift-diffusion model approach. *J. Cogn.* **6**, 46 (2023).
119. Box, G.E. & Tiao, G.C. *Bayesian Inference in Statistical Analysis* (John Wiley & Sons, 2011).
120. Tibshirani, R. J. & Efron, B. An introduction to the bootstrap. *Monogr. Stat. Appl. Probab.* **57**, 1–436 (1993).
121. Casella, G. & Berger, R. *Statistical Inference* (Chapman and Hall/CRC, 2024).

## Acknowledgements

We thank Daphne Cornelisse for creating Figs. 1 and 2, as well as Martijn Wisse, Todd Hester, Ran Wei, and Trent Victor for providing helpful feedback on the manuscript.

## Author contributions

**J.F.S.:** Methodology, Software, Validation, Formal analysis, Investigation, Writing - Original Draft, Visualization. **J.E.:** Conceptualization, Validation, Writing - Original Draft, Supervision, Project administration, Funding acquisition. **L.J.:** Formal analysis, Writing - Review & Editing, Visualization. **M.O'.K.:** Writing - Review & Editing. **J.M.:** Writing - Review & Editing. **J.K.:** Writing - Review & Editing, Supervision. **A.Z.:** Resources, Writing - Review & Editing, Supervision, Project administration, Funding acquisition.

## Competing interests

J.E., M.O'.K., L.J., and J.M. were employed by Waymo LLC and conducted the research without any external funding from third parties. Techniques discussed in this paper may be described in U.S. Patent Application Nos. 18/614,428 and 63/657,623. Delft University of Technology received funding from Waymo LLC for parts of the research carried out by J.F.S., but J.F.S. received no direct financial benefit for his contributions to the paper. The remaining authors declare no competing interests.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s41467-026-73345-0>.

**Correspondence** and requests for materials should be addressed to Johan Engström.

**Peer review information** *Nature Communications* thanks Conor Heins, Franck Mars and the other anonymous reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026