



Pedestrian crossing decisions can be explained by bounded optimal decision-making under noisy visual perception

Yueyang Wang^{a,*}, Aravinda Ramakrishnan Srinivasan^{a,d}, Jussi P.P. Jokinen^b, Antti Oulasvirta^c, Gustav Markkula^a

^a Institute for Transport Studies, University of Leeds, Leeds, LS2 9JT, UK

^b Faculty of Information Technology, University of Jyväskylä, Jyväskylä, 40100, Finland

^c Department of Information and Communications Engineering, Aalto University, Espoo, 00076, Finland

^d Chair and Institute of Highway Engineering (ISAC), RWTH Aachen University, Aachen, 52074, Germany

ARTICLE INFO

Dataset link: <https://github.com/YYWang98/Pedestrian-crossing-decisions.git>

Keywords:

Pedestrian behaviour
Computational rationality
Noisy perception
Reinforcement learning

ABSTRACT

This paper presents a model of pedestrian crossing decisions based on the theory of computational rationality. It is assumed that crossing decisions are boundedly optimal, with bounds on optimality arising from human cognitive constraints. While previous models of pedestrian behaviour have been either 'black-box' machine learning models or mechanistic models with explicit assumptions about cognitive factors, we combine both approaches. Specifically, we mechanistically model noisy human visual perception and model reward considering human constraints in crossing, but we use reinforcement learning to learn boundedly optimal behaviour policy. The model reproduces a larger number of known empirical phenomena than previous models, in particular: (1) the effect of the time to arrival of an approaching vehicle on whether the pedestrian accepts the gap, the effect of the vehicle's speed on both (2) gap acceptance and (3) pedestrian timing of crossing in front of yielding vehicles, and (4) the effect on this crossing timing of the stopping distance of the yielding vehicle. Notably, our findings suggest that behaviours previously framed as 'biases' in decision-making, such as speed-dependent gap acceptance, might instead be a product of rational adaptation to the constraints of visual perception. Our approach also permits fitting the parameters of cognitive constraints and rewards per individual to better account for individual differences, achieving good quantitative alignment with experimental data. To conclude, by leveraging both RL and mechanistic modelling, our model offers novel insights into pedestrian behaviour and may provide a useful foundation for more accurate and scalable pedestrian models.

1. Introduction

1.1. Background

Pedestrian crossing behaviour is a key factor in the urban transport system, impacting traffic efficiency and safety (Leu et al., 2012). Developing accurate models for pedestrian crossing decisions is therefore essential for effective urban planning (González-Méndez et al., 2021; Pelorosso, 2020). In particular, with the rise of autonomous vehicles (AVs), understanding pedestrian behaviour

* Corresponding author.

E-mail addresses: mn20yw2@leeds.ac.uk (Y. Wang), aravinda@isac.rwth-aachen.de (A.R. Srinivasan), jussi.p.p.jokinen@jyu.fi (J.P.P. Jokinen), antti.oulasvirta@aalto.fi (A. Oulasvirta), G.Markkula@leeds.ac.uk (G. Markkula).

<https://doi.org/10.1016/j.trc.2024.104963>

Received 6 February 2024; Received in revised form 29 November 2024; Accepted 2 December 2024

Available online 11 December 2024

0968-090X/© 2024 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

is important for the safe and harmonious integration of AVs into the existing transport ecosystem (González-Méndez et al., 2021; Pelorosso, 2020; Duric et al., 2002; Banovic et al., 2019; Al-Shihabi and Mourant, 2001; Li et al., 2016). The unpredictable nature of pedestrian movements, influenced by a complex interplay of cognitive processes and environmental cues, presents a significant challenge in this regard (Camara et al., 2020; Crosato et al., 2023).

In pedestrian behaviour modelling, computational models are predominantly based on either mechanistic modelling or data-driven machine learning (ML) approaches. Mechanistic models, which are grounded in assumptions about underlying cognitive mechanisms, aim to accurately represent the underlying processes causing pedestrian behaviour (Fajen and Warren, 2003; Turnwald et al., 2016; Yang et al., 2020; Pekkanen et al., 2021). However, formulating these mechanisms in a way that generalises across a wide variety of situations and scenarios is highly non-trivial. Therefore, since real-world traffic situations are highly variable and complex, mechanistic models typically suffer from scalability or generalisability issues. On the other hand, data-driven ML models, which predict pedestrian movement by learning from large datasets of pedestrian trajectories, show promise in handling complex real-world situations (Dai et al., 2019; Abughalieh and Alawneh, 2020; Quan et al., 2021; Yin et al., 2021; Yuan et al., 2021; Zhang et al., 2022; Li et al., 2022). However, ML models come with inherent challenges related to interpretability, data dependence, and robustness in diverse conditions (Althoff and Lutz, 2018; Klischat et al., 2020). These issues will be discussed in more detail in Section 1.2.

Ideally, one would want a modelling approach combining the strengths of both mechanistic and ML models. One possible candidate for such an approach is available from *computational rationality* (Lewis et al., 2014; Gershman et al., 2015), sometimes also referred to as *resource rationality* (Lieder and Griffiths, 2020). Both computational rationality and resource rationality are theories about human cognition and behaviour, positing that humans behave *rationally*, i.e., *optimally* with respect to some utility or cost function, but that this optimality is *bounded* by the constraints imposed by the human cognition and body. For example, pedestrians in road traffic have perceptual, cognitive, and motor limits which constrain their behaviour. Rooted in the principle of expected utility maximisation, introduced by Von Neumann and Morgenstern (1947), and expanded by Herbert Simon's concept of bounded rationality (Simon, 1955), computational rationality applies these decision-making principles while considering human constraints. In recent years, the development of models based on these theories has become increasingly attainable due to advances in the field of deep reinforcement learning (RL), enabling the learning of optimal decisions through environmental interactions (Silver et al., 2016, 2017; Sutton and Barto, 2018). Modern deep RL allows us to find boundedly optimal behaviour policy across highly varied scenarios and learn interactive behaviours which – to the extent that the theory of computational rationality is correct – mirror human-like decision processes and behavioural adaptability (Jokinen et al., 2020, 2021; Chen et al., 2021a,b). Contrary to traditional mechanistic models, which often struggle with flexibility in adapting to diverse scenarios and integration of various principles, and ML models, which can lack interpretability and adaptability, the computational rationality framework, when combined with RL, has the potential to overcome these constraints (Oulasvirta et al., 2022).

There have been some initial studies on the use of computational rationality in driver behaviour modelling (Jokinen et al., 2021, 2022), but so far, none in the context of vulnerable road users. Here, we provide a first demonstration of the value of this modelling framework in the context of pedestrian behaviour. Unlike previous pedestrian behaviour models, which are either mechanistic models, suffering from scalability or generalisability issues due to the complexities in real-world traffic conditions, or data-driven machine learning models, with challenges related to interpretability and data dependence, computational rationality addresses these issues by explaining human behaviour through a causal link between subjective utility, capacities, experience, and observed actions (Howes et al., 2023), optimising utility under defined constraints. The main human constraint we are considering in this work is imperfect human visual perception, characterised by the noise inherent in the visual system's processing of dynamic stimuli. This constraint often leads to mistakes in how people estimate the distance and speed of oncoming traffic, which we model mechanistically based on work in cognitive neuroscience. Looming aversion, which refers to the instinctive tendency to avoid objects that appear to be rapidly growing in size as a cue for an impending collision (DeLucia, 2008; Tian et al., 2022), is also considered in our model. We test our model on data from a controlled virtual reality experiment on pedestrian road-crossing, and find that our model captures four phenomena observed in the human experiment, which have not been captured simultaneously by one model. Furthermore, our model results also show that speed-dependent gap acceptance – which has previously been described as due to biases in human TTA estimation (Petzoldt, 2014) – is a boundedly optimal behaviour given the particular noise characteristics of the human visual system. Moreover, we show how individual differences in road user behaviour can be efficiently modelled using the computational rationality approach by conditioning the RL on the human constraint parameters. This method, which has not been previously used in road user behaviour modelling, includes parameters which can vary between humans as inputs to the RL, to learn optimal policy across variations in these parameters.

This paper is organised as follows: The rest of this section describes empirical studies of the behavioural phenomena we wish to capture and previous research on pedestrian behaviour modelling. Section 2 describes the dataset we used and the modelling approach. The results of our proposed model are presented in Section 3, followed by Section 4, with a discussion of results and future research plans. Finally, Section 5 provides a conclusion.

1.2. Related work

Numerous empirical studies have investigated the key factors influencing pedestrian behaviour (Oxley et al., 2005; Jain et al., 2014; Sun et al., 2015; Asaithambi et al., 2016; Gorrini et al., 2018; Tian et al., 2022). We focus on four main empirical phenomena, some but not all of which have been captured by existing models:

Table 1
Key phenomena in pedestrian behaviour and corresponding literature.

Phenomenon	Petzoldt (2014) and Tian et al. (2022)	Pekkanen et al. (2021)	This paper
(1) TTA-dependent gap acceptance	✓	✓	✓
(2) Speed-dependent gap acceptance	✓	✓	✓
(3) Speed-dependent yielding acceptance			✓
(4) Stopping distance-dependent yielding acceptance		✓	✓

(1) *TTA-dependent gap acceptance*: A key behaviour of interest is *gap acceptance*, where pedestrians decide to cross a road based on the time or spatial distance between them and an approaching vehicle. One crucial factor affecting this decision is the *time to arrival (TTA)*. Researchers have shown that pedestrians are more inclined to accept gaps for crossing when the TTA is higher (Oxley et al., 2005; Lobjois and Cavallo, 2007; Petzoldt, 2014).

(2) *Speed-dependent gap acceptance*: Additionally, the dynamics of approaching vehicles play an important role (Schneemann and Gohl, 2016). Pedestrians are more likely to cross in a given TTA when oncoming vehicles travel at higher speeds, i.e., with a larger approach distance (Lobjois and Cavallo, 2007; Tian et al., 2022). This phenomenon has been described as a bias in the human estimation of TTA (Petzoldt, 2014; Sun et al., 2015). Tian et al. (2022) explained this speed-dependent crossing behaviour by visual looming, the perceived growth of an object's size as it approaches (DeLucia, 2008). Specifically, visual looming increases slowly at long distances, indicating that higher-speed vehicles might produce smaller collision threats to pedestrians for a given TTA, thus influencing pedestrians to feel safer when crossing in front of faster vehicles.

(3) *Speed-dependent yielding acceptance*: In scenarios where vehicles yield, the relationship between speed and pedestrian crossing behaviour tends to reverse. Pedestrians have been found to interpret lower vehicle speeds as an indication of yielding, especially when the vehicle is at a closer distance, thus increasing the probability of crossing (Tian et al., 2023). On the other hand, early crossing decisions (while the vehicle is still some distance away) are similar to the speed-dependent gap acceptance regardless of whether the car is yielding.

(4) *Stopping distance-dependent yielding acceptance*: Driver use of exaggerated deceleration, or 'short-stopping', serves as a cue for pedestrians to cross (Domeyer et al., 2019). Short-stopping enhances the pedestrian's perception of the driver's intent, increasing their confidence in crossing safely (Domeyer et al., 2019). Risto et al. (2017) argued that the stopping distance of a vehicle correlates with pedestrian willingness to cross in the yielding scenario. Notably, Tian et al. (2023) found that the car's braking behaviour mainly influences late crossing decisions. Specifically, pedestrians tend to cross more readily when a vehicle stops at a greater distance from the crosswalk, seemingly interpreting this as a clear indication of the driver's intent to yield.

In our study, we aim to model pedestrian behaviour and reproduce these four phenomena mentioned above, summarised in Table 1. These phenomena provide valuable insights and highlight some of the nuances of human decision-making in traffic interactions that computational models should capture to be useful for practical applications such as simulation environments and AV algorithms. As mentioned, there are many models which describe pedestrian behaviour in terms of hypothesised underlying mechanisms (Fajen and Warren, 2003; Giles et al., 2019; Pekkanen et al., 2021; Tian et al., 2022, 2023). For example, Petzoldt (2014) developed a logit-based gap acceptance model capturing speed-dependent gap acceptance, described as a bias due to a visual heuristic used by pedestrians. However, this heuristic is truly biased in the sense that it is suboptimal or whether it might be an optimal adaptation to human constraints. Later, Tian et al. (2022) proposed a gap acceptance model utilising the binary choice logit approach. Unlike previous models that relied on traffic gap cues, their model incorporated visual looming cues, resulting in an improved fit to observed data. Another modelling approach is based on the concept of evidence accumulation, describing decision-making as a noisy accumulation of evidence from a stimulus (Ratcliff et al., 2016), which has also been applied to pedestrian crossing behaviour modelling (Pekkanen et al., 2021). However, due to the complexity of this model, the authors fitted it with a single parameterisation across all participants and with limited quantitative goodness of fit to the experimental data. The complexity of this model also makes it difficult to extend it to more sophisticated scenarios. To extend the scope of modelled scenarios, Markkula et al. (2023) integrated a large number of existing psychological and cognitive theories, such as theories of sensory noise, Bayesian perception, evidence accumulation decision-making, and long-term valuation of action affordances. Although this model was capable of reproducing several empirically observed phenomena in human road user interaction, the authors highlighted the limitations of mechanistic modelling and underscored the need for cognitively and behaviourally informed ML. Overall, given the high complexity of human behaviour, relying on a single cognitive, mechanistic model seems insufficient to comprehensively describe pedestrian behaviour. The existing models' success at capturing the targeted phenomena is listed in Table 1. However, it is worth noting that the aim of our work is not only to capture these phenomena but also to propose a general framework for modelling pedestrian behaviour.

With advances in computing power and machine-learning methods, a growing number of researchers have adopted data-driven ML algorithms for modelling pedestrian behaviour. For instance, many models based on convolutional neural networks have been developed for pedestrian trajectory prediction due to their ability to process spatial inputs, such as images and video frames (Yi et al., 2016; Abughalieh and Alawneh, 2020; Kumamoto and Yamada, 2017; Doellinger et al., 2018). Moreover, recurrent neural networks have garnered attention given their efficacy in sequence prediction tasks, making them apt for pedestrian trajectory modelling (Alahi et al., 2016; Dai et al., 2019; Quan et al., 2021). Recently, the Transformer model, known for its ability to process sequences in parallel and capture long-range interactions, has also gained popularity in pedestrian behaviour modelling (Lorenzo et al., 2021; Yin

et al., 2021; Yuan et al., 2021). While the accuracy of these ML models can be impressive, they share several limitations. First, they often act as ‘black boxes’, meaning they may have the potential to predict behaviours but fail to reveal the reasons or mechanisms behind those behaviours. This intensive data-driven orientation can be challenging, especially when aiming for interpretability in AV systems, where understanding pedestrian intent is crucial (Srinivasan et al., 2023). Second, the lack of theoretical grounding of the behaviour can sometimes lead to overfitting, where the models might perform well on the training data but fail to generalise well to unseen data, potentially resulting in less reliable predictions in varied real-world scenarios (Xu et al., 2020). Third, their heavy reliance on large training datasets poses another challenge: collecting such extensive data under all possible road conditions is almost impossible, and critical scenarios are missing in most datasets (Diaz-Ruiz et al., 2022).

As mentioned in Section 1.1, computational rationality provides a potential middle ground, combining the best of mechanistic and ML models, but has only seen limited use in road user behaviour modelling. Jokinen et al. (2021) proposed a hierarchical RL model of multitasking behaviour in driving and fitted it to observed data (Jokinen et al., 2022). However, their focus was on in-vehicle multitasking rather than pedestrian–vehicle interactions.

1.3. Summary of research gaps

Existing models of pedestrian crossing behaviour fail to capture the four mentioned critical empirical phenomena simultaneously, and neglect essential human perceptual and cognitive constraints like noisy perception and looming aversion. Furthermore, no existing model combines mechanistic and machine learning approaches, or accounts for individual differences in pedestrian behaviour.

Here, we aim to quantitatively model the crossing decisions of human pedestrians when faced with constant-speed and yielding vehicles. There are three research objectives in this study:

(1) To develop a boundedly optimal model, integrating human constraints, with a focus on noisy perception and looming aversion, that is able to capture the four targeted empirical phenomena in human pedestrian crossing behaviours in scenarios where vehicles may or may not yield to the pedestrian.

(2) To fit the boundedly optimal model quantitatively to human crossing decision data on a per-individual level (i.e., with different model parameter values for different individuals).

(3) To assess how the human constraints and preferences integrated into the model influence the crossing decisions.

An early version of this work has been presented as a conference paper (Wang et al., 2023). This conference paper qualitatively explored the effect of only noisy visual perception on crossing decisions and only in scenarios with approaching vehicles at constant speed; in other words, it only partially addressed the first objective above, and not at all the second or third objectives.

2. Methods

2.1. Dataset

This study utilised a dataset sourced from a previous experiment reported by Giles et al. (2019). A visualisation of the experimental setup is provided in Fig. 1, which offers a bird’s-eye perspective. Twenty participants (age 24–60, average 27.9 years; 11 males and 9 females) took part in the study and were recruited from a university participant pool. We did not record the racial or ethnic identities of our participants in this study. Regarding the age influence, it is important to note that the experimental task only required button press responses, such that participant variation in motor ability was less of a concern than it would have been in a task requiring actual walking. Participants wore an HTC Vive Virtual Reality (VR) headset as part of the setup, immersing them in a virtual crossing task. The rendered VR space included a straight two-lane road, spanning a width of 5.85 m, and a zebra crossing at the participant’s initial location.

At the start of each trial, participants were positioned before the zebra crossing, facing straight across it. They had been instructed that they should look to the right for any oncoming traffic when they felt ready to begin the trial. This head movement triggered the start of the trial, with an oncoming vehicle initialised at an initial distance d_0 and speed v_0 . The experiment included a mix of scenarios, namely constant-speed and yielding scenarios. The detail of these scenarios is shown in Table 2, with the initial TTA $\tau_0 = d_0/v_0$ also listed. In six constant-speed scenarios, the vehicle appeared at a distance d_0 (all distances measured longitudinally along the road from the participant’s location to the front of the car) and maintained a constant-speed v_0 while approaching and passing the zebra crossing. In eight yielding scenarios, the vehicle appeared at initial distance and speed v_0 and d_0 , and immediately decelerated at a constant rate to stop at a distance d_{stop} from the participant. Participants were instructed to press the HTC Vive’s controller button when they felt safe to cross. It should be noted that the vehicle’s behaviour was not influenced by the pedestrian’s decision. This setup required participants to assess crossing safety independently. Once the button was pressed, the system recorded the *Crossing Initiation Time (CIT)*, defined as the interval between the start of the trial and the moment when the button press occurred. At this point, the position of the participant’s point of view in the virtual environment (i.e., the VR ‘camera’) moved across the zebra crossing at the speed of 1.31 m/s. Each participant faced 6 unique constant-speed trials and 8 unique yielding trials in a randomised order, aggregating to a total of 280 trials that were used for the model’s validation.

The low number of repetitions per participant was adopted by Giles et al. (2019) to limit behavioural adaptation effects. The resulting dataset is relatively small in size. However, since all of the main phenomena we target have also been reported in other experiments, the sample size is not a main concern here. We adopted this dataset here because of the button-press paradigm, which provides crossing onset distributions with minimal impact from motor variability, aligning with our present modelling emphasis on perception rather than motor control.

Table 2

Detailed description of vehicle approach scenarios in the experiment. The table lists each scenario type along with key parameters: Initial Vehicle Speed (v_0 in m/s), Initial Distance (d_0 in m) from the pedestrian, Initial TTA (τ_0 in s), and Stopping Distance (d_{stop} in m) for Yielding scenarios. ‘N/A’ indicates not applicable for constant speed scenarios.

Scenario type	v_0 (m/s)	d_0 (m)	τ_0 (s)	d_{stop} (m)
Constant speed	6.94	15.90	2.29	N/A
	13.89	31.81	2.29	N/A
	6.94	31.81	4.58	N/A
	13.89	63.61	4.58	N/A
	6.94	47.71	6.87	N/A
	13.89	95.42	6.87	N/A
Yielding	6.94	15.90	2.29	4
	13.89	31.81	2.29	4
	13.89	31.81	2.29	8
	6.94	31.81	4.58	4
	13.89	63.61	4.58	4
	13.89	63.61	4.58	8
	6.94	47.71	6.87	4
	13.89	95.42	6.87	4

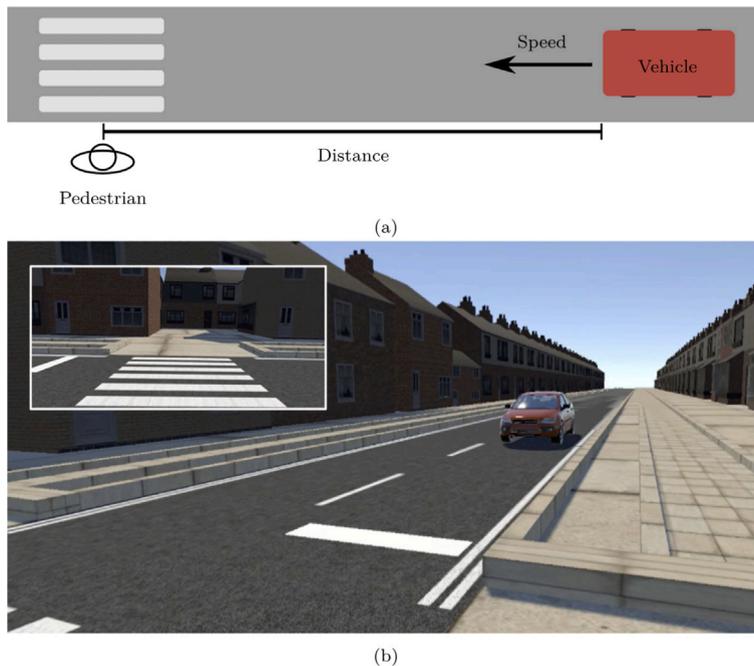


Fig. 1. (a) Birds-eye view of the experiment. (b) A sample view of the virtual scene, as shown by the head-mounted display, at the beginning of each trial (inset) and as participants turned their heads to look for oncoming traffic.

Source: Pekkanen et al. (2021). Variable-drift diffusion models of pedestrian road-crossing decisions. *Computational Brain & Behaviour*, 1–21. This image is available under a Creative Commons Attribution 4.0 International License.

2.2. Model variants

In this section, we explain our two main mechanistic assumptions, rooted in cognitive science and neuroscience, and introduce different model variants derived from these assumptions:

- **Assumption of noisy perception:** We assume that the agent’s perception of the environment is inherently noisy and imperfect. This noisy observation obtained by the agent is according to the principle of the human visual system, i.e., the sensory input received by our human visual system is noisy (Faisal et al., 2008).
- **Assumption of looming aversion:** Our second assumption involves the natural aversion to looming objects (DeLucia, 2008; Tian et al., 2022). Looming aversion refers to the instinctive tendency to avoid objects that appear to be rapidly growing in size, as this is often a cue for an impending collision (DeLucia, 2008; Tian et al., 2022).

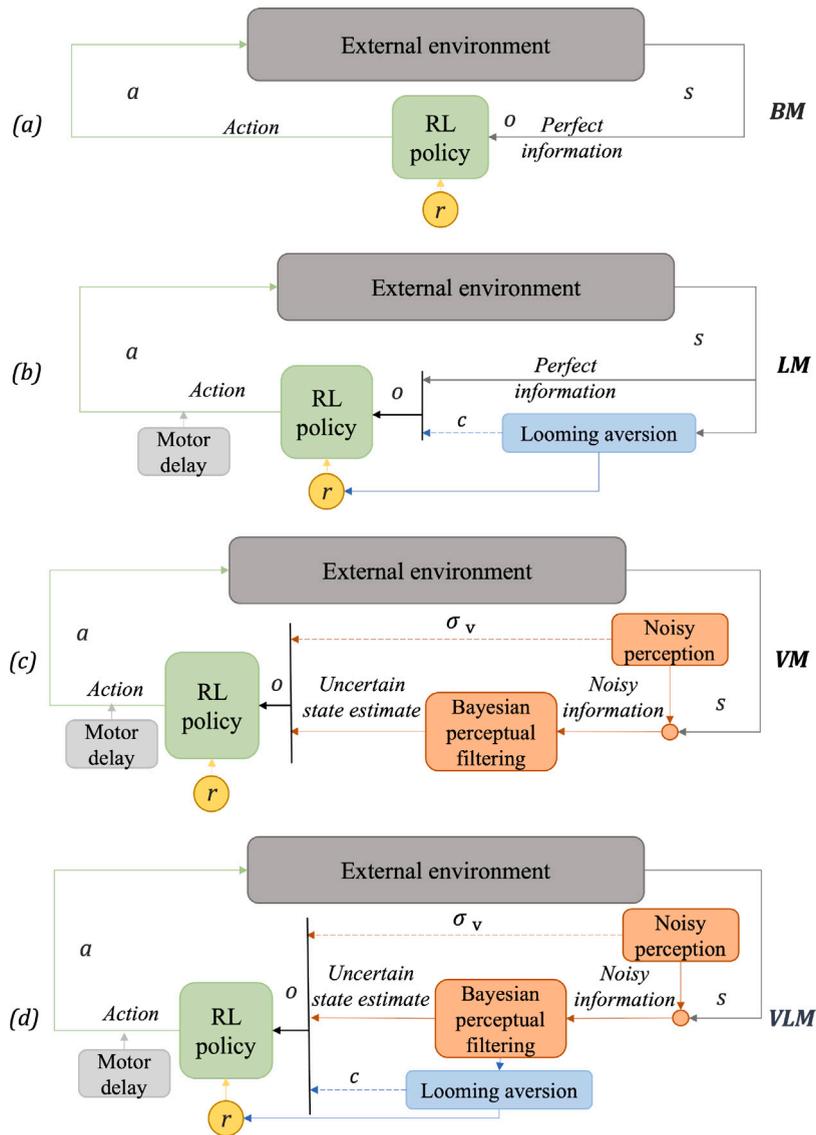


Fig. 2. Comparison of models. (a) BM: Baseline Model. (b) LM: Looming aversion only. (c) VM: Visual constraints only. (d) VLM: Visual constraints and looming aversion. σ_v and c represent the sensory noise and looming aversion weight respectively. s , o , a , and r represent the state, observation, action and reward respectively.

These two assumptions guide the development of our model variants. Each variant incorporates the assumptions in different ways to explore how they individually and collectively influence pedestrian behaviour. The following subsection outlines the four main model variants derived from these assumptions, as visualised in Fig. 2:

(1) **Baseline Model (BM)**: The BM serves as our control variant, assuming an ideal observer with neither visual constraints nor looming aversion. This model establishes a baseline against which to measure the impact of our two key assumptions.

(2) **Looming Model (LM)**: The LM considers the impact of looming aversion on crossing decisions. It allows us to understand how this aversion, independent of visual constraints, can affect pedestrian behaviour.

(3) **Visual constraints Model (VM)**: Conversely, the VM examines the role of visual constraints alone. It helps us investigate how the noisy sensory information influences pedestrian crossing.

(4) **Visual constraints and Looming Model (VLM)**: The VLM combines both visual constraints and looming aversion to provide insights into their joint effect.

Each combination of model assumptions mentioned above defines an RL problem. In the next section, we describe this RL problem in detail for the different model variants.

2.3. Reinforcement learning problem

Our model based on computational rationality used RL to derive the near-optimal behaviour under constraints. In RL, the agent chooses the action by following a policy π , which yields a probability $\pi(s, a) = p(a | s)$ of taking a particular action from the given state. The optimal policy π^* , maximises the value function:

$$V^*(s) = \max_a \left[R(s, a) + \gamma \sum_{s' \in S} T(s, a, s') V^*(s') \right] \quad (1)$$

where γ is the discount factor, a value between 0 and 1, which balances the importance of immediate rewards versus future ones. Here, $V^*(s)$ represents the optimal value function at state s , and $V^*(s')$ represents the optimal value at state s' .

In our study, we assumed that humans make decisions under uncertainties from the noisy perception system, such that the overall task can be understood as a Partially Observable Markov Decision Process (POMDP). In a POMDP, the agent does not have direct access to the actual state S ; instead, it receives observations that may only partially or noisily reflect the actual state. The POMDP is represented by a tuple $\langle S, A, T, R, O \rangle$, where S is a set of states, A is a set of actions, T is the transition function, R is the reward function, and O is the set of possible observations received by the agent.

State space S . At each time step t , the environment is in a state $s_t \in S$. This is the positions and velocities of the pedestrian (the RL agent) and the vehicle. All model variants in our study share the same state space, encapsulating these five critical variables: the pedestrian's position x_p, y_p , the vehicle's position x_{veh}, y_{veh} , the vehicle's velocity v , and time step t . The simulation state update time step is set to 0.1 s, a duration suitable for our dataset and model.

Action space A . At every time step t , the agent executes an action a_t from the set A . For the purposes of this study, and consistent with the experimental button press, the agent's choices are binary: either to 'Go' or to 'Not Go'. Upon selecting 'Go', the agent proceeds straight at the speed of 1.31 m/s, as in the experiment. Recognising that human reaction times vary, we introduced a motor delay to simulate this aspect. This delay, implemented after the 'Go' decision, was sampled from a Gaussian distribution with a mean of 0.6 s and a standard deviation of 0.2 s. After the decision was made, the simulation progressed until the agent either crossed the road safely or a collision happened.

Transition T . The transition function defines how the current state s_t changes to the next state s_{t+1} based on the action a_t . In our model, when the 'Not Go' action is selected, the vehicle's movement follows kinematic equations with the given speed, and the agent's position remains unchanged. Conversely, when the 'Go' action is chosen, whether the collision happens is calculated, and the corresponding reward is given to the agent. Then, the simulation finishes.

Reward R . In the experiment where the dataset was collected, the participant's task was to cross the road as soon as they felt safe to do so, either before or after the car had passed them (Giles et al., 2019). Therefore, we assume that the participants in this experiment (presumably similar to pedestrians in real traffic) wanted to cross safely but with minimal loss of time. We designed the reward structure accordingly: The agent will be given a reward of 20 when crossing the road without collision and a reward of -20 if a collision happens. These reward parameters were shared across model variants. After some initial manual testing, they were fixed to yield reasonable crossing behaviour of the BM agent but before structured fitting of any other model parameters to the human data. A time penalty, a negative reward of $0.01 \times t$, will also be given to the agent when the episode terminates.

Regarding the visual looming, this phenomenon is represented mathematically as inverse τ —the ratio of a vehicle's optical expansion rate to its size on the observer's retina, which serves as an estimate of the inverse TTA (DeLucia, 2015; Markkula et al., 2016). We incorporated this concept into the reward function of the LM and VLM models, as shown in panels (b) and (d) of Fig. 2, to account for the looming aversion in pedestrian decision-making.

The reward function r is defined as follows:

$$r = \begin{cases} \max(-20, \min(+20, 20 - 0.01 \cdot t - c \cdot \frac{1}{\hat{\tau}})), & \text{if arrival} \\ -20, & \text{if collision} \end{cases} \quad (2)$$

where c is the weight of the looming aversion. For model variants with non-noisy perception, $\tau = \frac{x_{veh} - x_p}{v}$, and for those with noisy perception, $\hat{\tau} = \frac{\hat{x}_{veh} - \hat{x}_p}{\hat{v}}$, where \hat{x}_{veh} and \hat{v} are noisy estimates of vehicle position and speed (see further below). The reward r is bounded within the range $[-20, +20]$ to prevent extreme values from influencing the model training. As the focus of this study is on the potential effect of noisy perception and the looming aversion to the crossing decision, we have kept the reward function simple; future work can further refine it to better capture human preferences.

Observation space O . The agent receives observation $o_t \in O$ at each time step as shown in Table 3. We tested different formulations of the observation space: In BM—our simplest, baseline model as shown in panel (a) of Fig. 2—the exact position and velocity of both the vehicle and the pedestrian are presented as inputs to the agent. Conversely, in the models with visual constraints assumption, VM and VLM, illustrated in panels (c) and (d) of Fig. 2, the position and velocity details provided to the agent are subject to noise, emulating the inherent uncertainty and imperfection in human visual perception, with an angular noise introduced at the level of the human retina (Kwon et al., 2015). We assumed that the agent observes the position of the other agent along its line of travel by observing the angle below the horizon of the other agent (Ooi et al., 2001; Markkula et al., 2023), with a constant Gaussian

Table 3

Observation space variables for each model variant. ‘✓’ indicates the variable is directly observed. Here, x_p, y_p denote the pedestrian’s position, x_{veh}, y_{veh} the vehicle’s position, v the vehicle’s velocity, \hat{x}_{veh}, \hat{v} the Kalman filter’s estimates of the vehicle’s position and velocity, P_p and P_v are the variances in the Kalman filter’s estimate of position and speed of the approaching vehicle, σ_v represents the standard deviation of sensory noise, c the looming aversion weight, and t the time step.

Model	x_p	y_p	x_{veh}	y_{veh}	v	Additional variables
BM	✓	✓	✓	✓	✓	t
LM	✓	✓	✓	✓	✓	c, t
VM	✓	✓	Est. (\hat{x}_{veh})	Est. (\hat{y}_{veh})	Est. (\hat{v})	P_p, P_v, σ_v, t
VLM	✓	✓	Est. (\hat{x}_{veh})	Est. (\hat{y}_{veh})	Est. (\hat{v})	P_p, P_v, σ_v, c, t

angular noise of standard deviation, σ_v , which could vary between pedestrians. In practice, this means that the pedestrian observes the vehicle’s distance along the road with a distance-dependent noise of standard deviation (Markkula et al., 2023):

$$\sigma_x(t) = |d_1(t)| \left(1 - \frac{h}{d(t) \cdot \tan(\arctan \frac{h}{d(t)} + \sigma_v)} \right), \quad (3)$$

where $d_1(t)$ is the longitudinal distance between the vehicle and the crossing point, $d(t)$ is the distance between the agent and the approaching vehicle, and h is the eye height over the ground of the ego agent, which is set to 1.6 m for all pedestrian agents for simplicity.

Additionally, there is evidence that the human perception system interprets its noisy input in a Bayes-optimal manner, and Bayesian methods have been successful in modelling perception and sensorimotor control (Kwon et al., 2015; Knill and Pouget, 2004; Stocker and Simoncelli, 2006). Therefore, we used a Kalman filter as a model of the human visual perception to perceive the environment (Kwon et al., 2015; Markkula et al., 2023). In the initialisation of the Kalman filter, we used distinct prior distributions for position and velocity. Each distribution is centred on the actual vehicle state, with its standard deviation set equal to that of the initial set of values. At each time step, the Kalman filter receives noisy positional data regarding the other agent (vehicle). The filter then produces estimates of the vehicle’s position and velocity, along with their respective uncertainties. These filtered estimates represent the agent’s belief state, reflecting the Bayes-optimal inference of the vehicle’s state based on the noisy observations. These belief estimates are then fed into the RL policy as inputs, together with the precise self-position and velocity of the agent itself and the visual noise parameter σ_v .

Furthermore, the LM and VLM models incorporate the concept of visual looming. This looming aversion, modelled through inverse τ as shown in Eq. (2), affects the agent’s perception of potential collision risks. We gave the weight for the looming aversion, c , as input to the RL policy.

To distinguish σ_v and c from the parameters of the policy neural network (connection weights and biases), we will refer to these two parameters as non-policy parameters. By feeding the non-policy parameters as inputs to the RL policy, we are not implying that the human agent ‘observes’ its own parameter values. Instead, we are just conditioning the RL policy on these non-policy parameters, as a more convenient alternative to learning entirely different policies for different parameter combinations.

2.4. Reinforcement learning algorithm

Deep Q-Networks (DQNs) optimise the Q-function to estimate the expected utility of actions in given states, balancing immediate and future rewards. The Q-function is defined as:

$$Q(s, a) = r + \gamma \max_{a'} Q(s', a'), \quad (4)$$

where s represents the current state, a the action taken in state s , and r the immediate reward received after taking action a in state s . The DQN algorithm is suitable for problems with a continuous state space and a discrete action space, such as in our case. We adopted Duelling DQN, which enhances standard DQN by separating the action selection from evaluation with its unique structure of a value function $V(s)$ and an action advantage function $A(s, a)$. This allows for more precise evaluations in scenarios where actions have similar outcomes, improving decision-making accuracy by focusing on both state value and action advantages (Van Hasselt et al., 2016; Wang et al., 2016).

We represented the $Q(s, a)$ values using a fully connected feedforward neural network with two hidden layers of 512 and 256 nodes. It has been observed that larger networks tend to yield more stable policies, and policies are more robust to noise (Xie et al., 2019). The learning rate and discount factor were set to 0.0001 and 0.99, respectively. To encourage exploration, we implemented an ϵ -greedy algorithm: at each time step t , the system either randomly selects an action with probability ϵ or chooses the action with the highest Q value with probability $1 - \epsilon$ (Wunder et al., 2010). We initialised ϵ at 1 and decreased it by 5^{-5} in each learning step. The minimum value of ϵ was set to 0.001. We trained model M, VM and LM over 25,000 episodes for convergence. Convergence here means that rewards stabilise and no longer increase over time. The choice of this criterion is rooted in the understanding that the model has likely learned an optimal or near-optimal policy once the rewards stop improving significantly. For the model VLM, we extended the training to 45,000 episodes due to its added complexities from both visual constraints and looming aversion assumptions.

To avoid the agent learning a simplistic strategy of always crossing immediately, we also included scenarios with an initial TTA of 1 s, in which safe crossing was only feasible after the vehicle’s arrival.

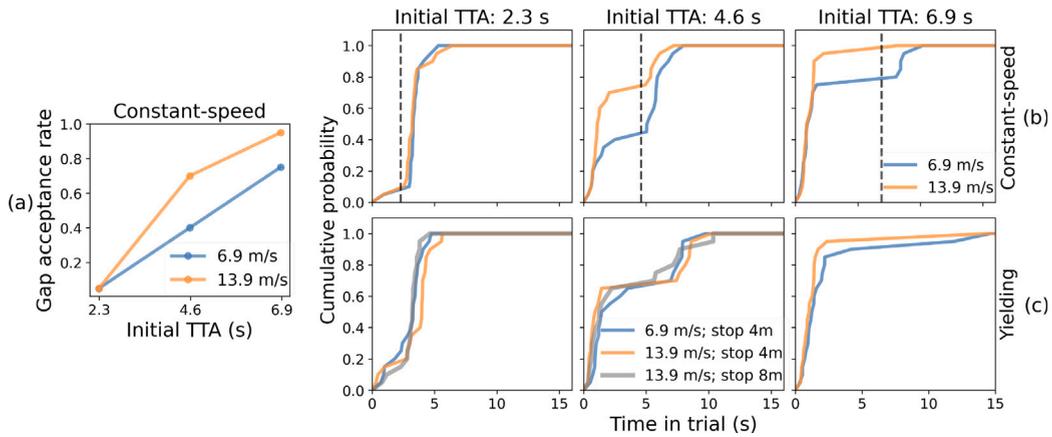


Fig. 3. Empirical results. (a) Gap acceptance rate in constant-speed scenarios. (b) Cumulative probability of Crossing Initiation Time (CIT) in constant-speed scenarios; black dashed vertical lines indicate the times vehicles passed pedestrians. The X -axis means the time elapsed from the beginning of each trial. (c) Cumulative probability of CIT in yielding scenarios. Note: In yielding scenarios, there are only two conditions in the initial TTA of 6.9 s as shown in Table 2, so there is no grey line in the third panel of yielding scenarios.

2.5. Fitting of non-policy parameters

Our most complex model variant, VLM, has two free model parameters, σ_v and c . Testing different values for these non-policy parameters is essential because they alter the environment in the RL problem, thus influencing the agent's crossing decision.

Learning a separate RL policy for each possible combination of non-policy parameter values is computationally expensive (Howes et al., 2023; Li et al., 2023). Recent studies have adopted a more efficient approach within the computational rationality framework to address this challenge. This method involves conditioning the RL policy on non-policy parameters by integrating them as additional inputs to the model (Keurulainen et al., 2023). In this case, we provided σ_v and c as additional inputs to the RL policy during learning, as illustrated by the dashed line in Fig. 2(d). Utilising this approach, the model can learn a boundedly optimal policy for arbitrary values of the non-policy parameters.

We do not know the correct values for the noise magnitude parameter, σ_v , and the weight for the looming aversion, c , during the training phase. Moreover, these values might differ for individual participants in the experiment. Therefore, during RL training, for each episode we sampled values for these parameters from a range of σ_v values (from 0 to 1 in increments of 0.1) and c values (from 0 to 100 in increments of 10). Then, after RL training was complete, we tested this learned RL policy across the entire 10×10 grid of non-policy parameters to make predictions about model crossing decisions in all fourteen experiment scenarios for each parameter value combination.

Building on this exhaustive parameter search, we determined the σ_v and c that fitted the experimental data best for each participant by likelihood maximisation. We estimated the probability density function of CIT predicted for each model by kernel density estimation separately for each of the fourteen scenarios (Table 2). This allowed us to calculate the model likelihood of each σ_v and c for each participant as the product of multiplying the model-predicted probability density at the participant's observed button press times.

3. Results

3.1. Empirical results

Panel (a) of Fig. 3 shows that the pedestrian gap acceptance rate increases with the initial TTA ($p < 0.05$ and $p < 0.05$ for 6.9 m/s and 13.9 m/s respectively; Chi-squared test), which aligns with previous research (Oxley et al., 2005; Lobjois and Cavallo, 2007; Petzoldt, 2014). From Fig. 3, we can also see a speed-dependent gap acceptance rate, suggesting that pedestrians were more likely to accept the gap as the vehicle's speed increased, given the same initial TTA, again aligning with findings from previous studies (Lobjois and Cavallo, 2007; Tian et al., 2022). However, while this speed effect was particularly evident at initial TTA of 4.6 s and 6.9 s, the associated p -values (0.06 and 0.08) are not statistically significant ($p < 0.05$), possibly due to our limited sample size.

As shown in panel (b) of Fig. 3, this speed-dependency is also visible in the CIT metric (Tian et al., 2022), which in our case denotes the time from the start of each trial to the pedestrian's crossing decision. This phenomenon was most evident in the initial TTA of 4.6 s ($p < 0.05$; Wilcoxon Signed-Rank test).

As for the yielding scenarios, where the car stopped to let the participant cross, it can be seen in Fig. 3(c) that the car speed's impact on CIT was most pronounced under a higher initial TTA condition of 6.9 s ($p < 0.05$; Wilcoxon Signed-Rank test). Specifically, the third graph in panel (c) of Fig. 3 shows the orange curve ascending more rapidly than the blue curve, which means that more

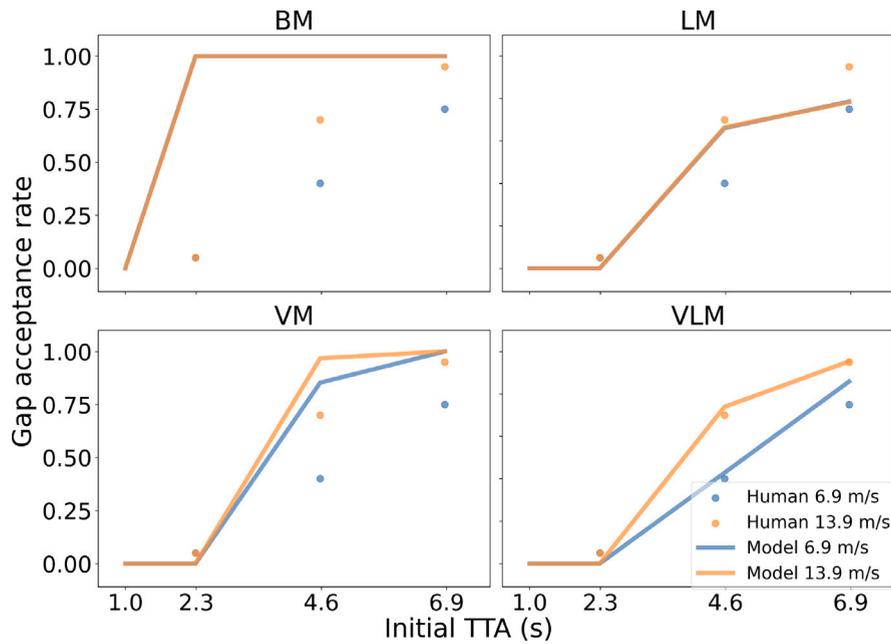


Fig. 4. Gap acceptance rate by human participants and different models. BM: Baseline model. LM: Model with the looming aversion assumption. VM: Model with the visual constraints assumption. VLM: Model with both visual constraints and looming aversion assumptions. Note: the lines of the model LM and BM overlap each other.

pedestrians cross earlier when the vehicle's initial speed is higher at higher TTA conditions. This speed-dependent trend is the same as in constant-speed scenarios. Conversely, an inverse pattern emerged for a lower initial TTA of 2.3 s. More pedestrians chose to cross earlier when the approaching vehicle's initial speed was low during yielding ($p < 0.05$; Wilcoxon Signed-Rank test), indicating a speed-dependent yielding acceptance. This speed-dependent yielding behaviour is illustrated by the blue curve's quicker increase compared to the orange curve in the first graph of the panel (c) in Fig. 3. This finding is in line with the work by Ackermann et al. (2019) and Tian et al. (2023), suggesting that pedestrians tend to interpret low speed in itself as indicative of vehicle yielding. Interestingly, a combination of these two phenomena is observed at the initial TTA of 4.6 s, with a speed-dependent gap acceptance in the early phase of the scenario and a speed-dependent yielding acceptance behaviour emerging in the late phase of the trial. Therefore, across the entire scenario, the mean CIT is not significantly affected by car speed ($p = 0.21$; Wilcoxon Signed-Rank test).

Finally, stopping distance-dependent yield acceptance can also be seen in Fig. 3. More pedestrians tended to cross early when the vehicle decelerated at a higher rate, as shown by the orange line lagging behind the grey line in the yield acceptance phase of the scenario (from about 2 s in the TTA 2.3 s scenario and about 5 s in the TTA 4.6 s scenario), indicating that greater vehicle deceleration encourages early crossing (Ackermann et al., 2019). This effect of stopping distance on CIT was statistically significant at the initial TTA of 2.3 s ($p < 0.05$; Wilcoxon Signed-Rank test).

In sum, the empirical observations show indications of the four main phenomena we are targeting, even if some are present as trends rather than as statistically significant effects in this relatively small dataset.

3.2. Model results

After observing these key behaviours in empirical settings, we now investigate whether our model reproduces these phenomena. All model variants converged; a plot of the reward curves can be found in Appendix A.1.

3.2.1. Evaluation of model outcomes against observed phenomena

We evaluated our models based on their ability to replicate our four targeted key pedestrian behaviours identified in empirical studies.

First, we analysed the results of the baseline model (BM). The results of BM for the constant-speed scenarios are shown in the upper-left panel of Fig. 4 and the second row in Fig. 5. As shown in Fig. 4, the BM agent decided not to cross at the initial TTA of 1 s but consistently accepted all safe crossing opportunities available during initial TTA conditions 2.3 s, 4.6 s and 6.9 s, where it crossed at the very start of the trial (shown in Figs. 5 and 6). This behaviour can be attributed to the agent's perfect information about the vehicle's state, permitting it to achieve optimal behaviour for an agent whose goal is to cross the road safely in the least possible time. The BM model variant does not capture any of the four targeted phenomena. Below, we will consider these phenomena one at a time.

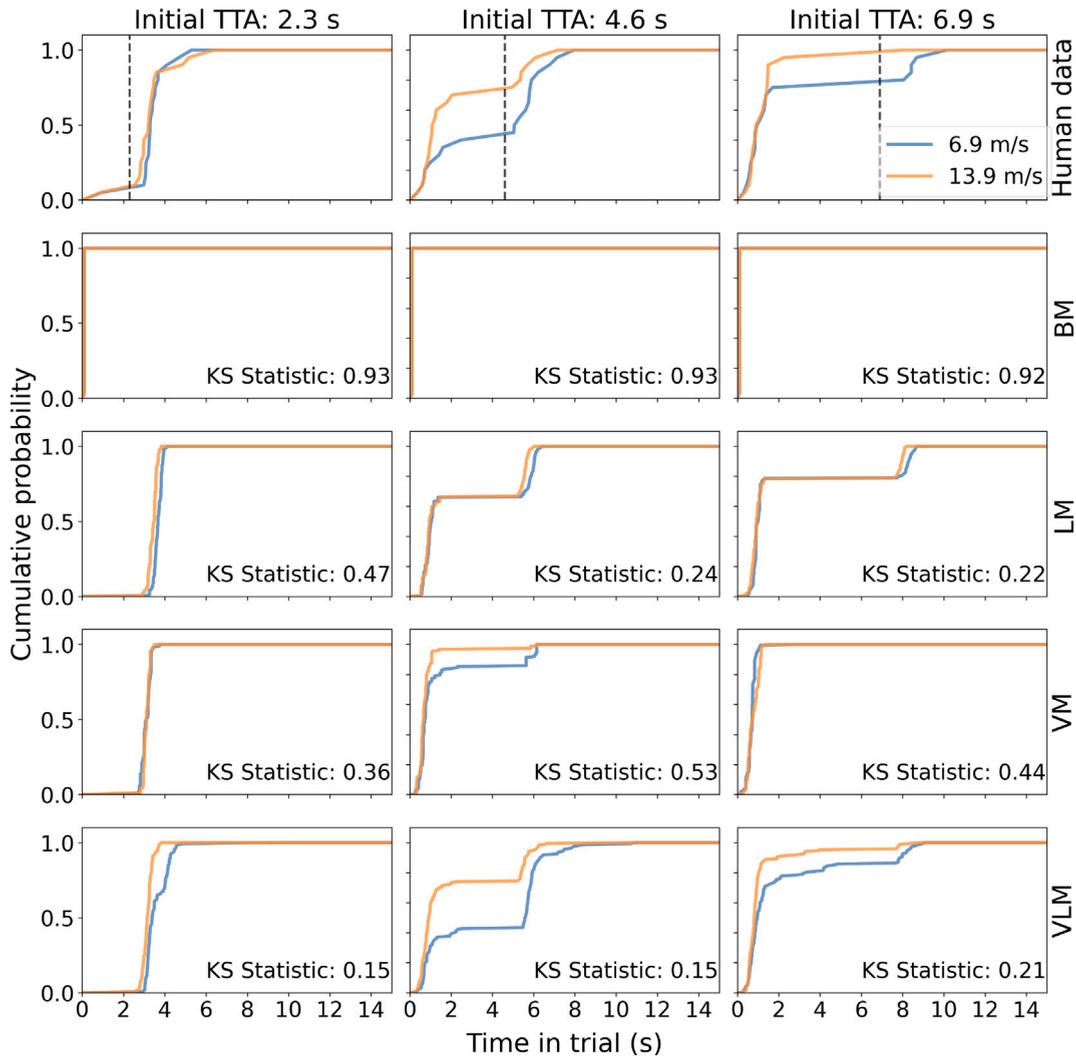


Fig. 5. Cumulative probability for Crossing Initiation Time (CIT) in constant-speed scenarios. Black dashed vertical lines in the human data indicate the times vehicles passed pedestrians. Vehicles passed pedestrians at the same time in the trial for each initial TTA condition. The X-axis represents the time elapsed from the beginning of each trial. The KS statistic quantifies the maximum divergence between the cumulative distribution functions of human data and model results, indicating their distributional similarity. Model abbreviations are as defined in Fig. 4.

TTA-dependent gap acceptance. The VM, LM, and VLM models all qualitatively captured the TTA-dependent gap acceptance phenomenon, showing increasing gap acceptance rates with rising initial TTA conditions. This pattern is particularly evident in the LM model, as seen in the upper-right panel of Fig. 4. The VLM model, integrating both visual constraints and looming aversion (shown in the bottom-right panel of Fig. 4), most accurately reflected the human data, demonstrating the most human-like sensitivity to initial TTA conditions.

Speed-dependent gap acceptance. While including just looming aversion (model LM) was enough to capture the effect of TTA on gap acceptance, this model did not show any effect of vehicle speed on gap acceptance. This can be seen in Fig. 4, where the LM lines for two speed conditions overlap entirely. On the other hand, the inclusion of visual constraints (model VM) caused the model to exhibit speed-dependent gap acceptance behaviour by separating the gap acceptance curves at the initial TTA of 4.6 s. The VLM model again exhibited the best similarity with human data, showing a clear speed dependency in the gap acceptance rate.

Speed-dependent yielding acceptance. As shown in the bottom row in Fig. 6, at the initial TTA condition of 2.3 s, speed-dependent yielding acceptance emerged in the results of VLM, similar to the human data, visible as a quicker increase of the blue curve than the orange curve. In other words, the agent was more likely to cross the road earlier when the initial speed of the approaching vehicle was lower. Furthermore, and again in line with the human data, at the initial TTA condition of 6.9 s, the initial speed of the vehicle has the opposite effect on CIT (speed-dependent gap acceptance), and in the 4.6 s initial TTA condition both of the two

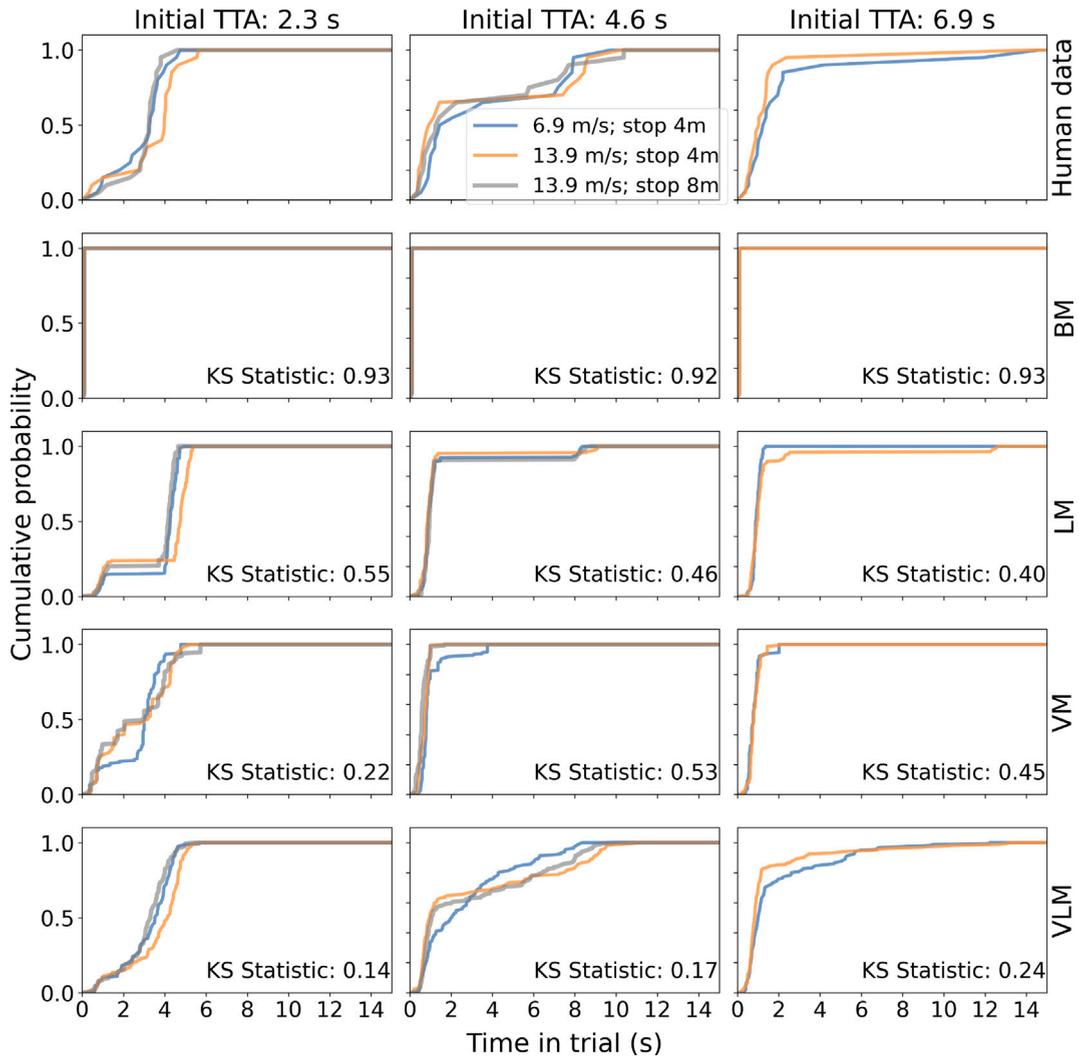


Fig. 6. Cumulative probability for CIT in yielding scenarios. See Fig. 5 for further details.

Table 4
Quantitative assessment of model performance.

Model type	Log Lik.	Params.	AIC	MAD (s)
BM	-1289	0	2578	2.90
LM	-594	20	1227	0.94
VM	-588	20	1215	1.30
VLM	-533	40	1146	0.34

Note: Log Lik. refers to Log Likelihood, Params. refers to Free Parameters, and MAD refers to the mean absolute deviation across all scenarios.

phenomena are visible. These speed effects can, to some extent, be observed in the results of LM and VM at the initial TTA of 2.3 s as well, but not for the 4.6 s and 6.9 s initial TTA conditions.

Stopping distance-dependent yielding acceptance. The VLM agent replicated this phenomenon both at the initial TTA condition of 2.3 s and 4.6 s, i.e., the agent had a greater tendency to cross earlier when the vehicle stopped at a greater distance, visible as the orange line lagging the grey line in the CIT plots. Notably, at the initial TTA of 4.6 s, this tendency appeared only in the late crossing decisions (from about 5 s). This is probably because the deceleration information has a stronger effect on the crossing decision when the vehicle is not far from the pedestrian, in line with Tian et al. (2023). The LM model also captured this pattern, especially at the initial TTA condition of 2.3 s. However, this pattern did not appear in the results of the VM model.

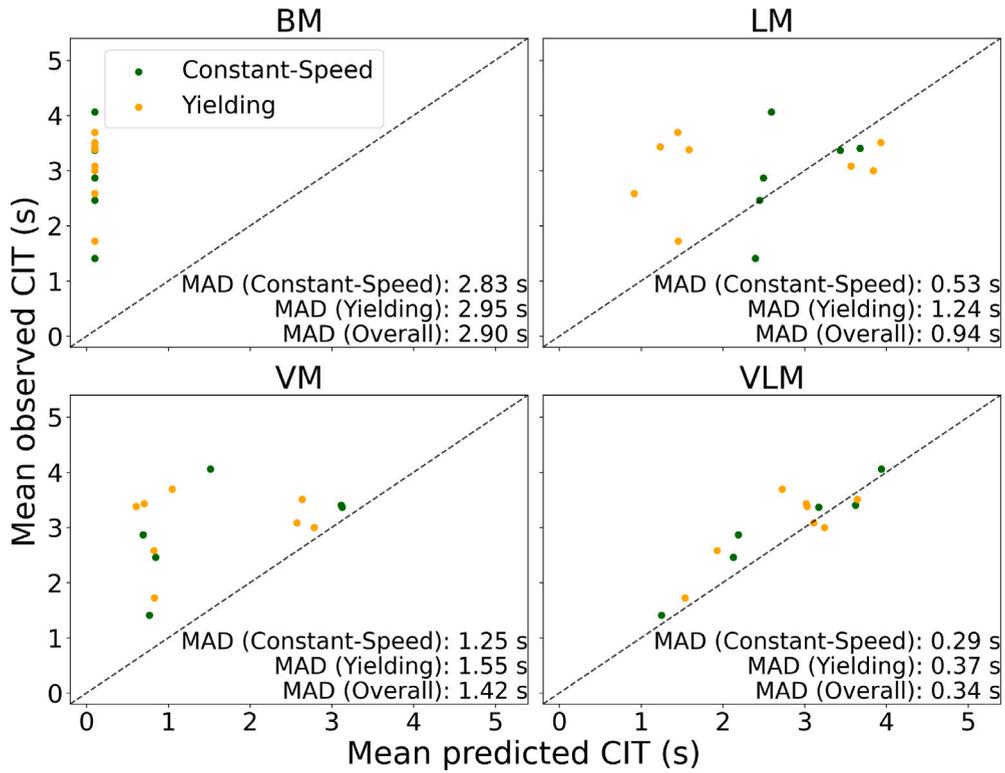


Fig. 7. Predicted vs observed mean Crossing Initiation Time (CIT) across the different scenarios. This comparison illustrates the model’s performance in estimating CIT in relation to actual measurements, with the dotted line representing an ideal prediction where observed and predicted values match perfectly.

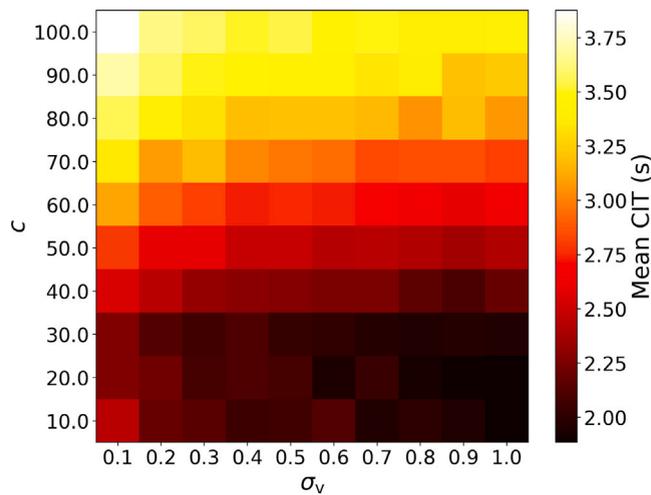


Fig. 8. Heatmap of mean CIT for different values of the two non-policy parameters. Different colours represent the average CIT values across all different scenarios.

3.2.2. Quantitative assessment of model performance

We assessed model performance using log-likelihood, Akaike Information Criterion (AIC), and Mean Absolute Deviation (MAD). Higher log-likelihood and lower AIC values indicate better model fit and simplicity, respectively. At the same time, a smaller MAD suggests a closer match between predicted and observed crossing times. As shown in Table 4 and Fig. 7, the VLM model demonstrated superior performance across all metrics: it achieved the highest log-likelihood (−533), the lowest AIC (1146), and the smallest MAD (0.34 s), indicating it most accurately represents real-world pedestrian behaviour.

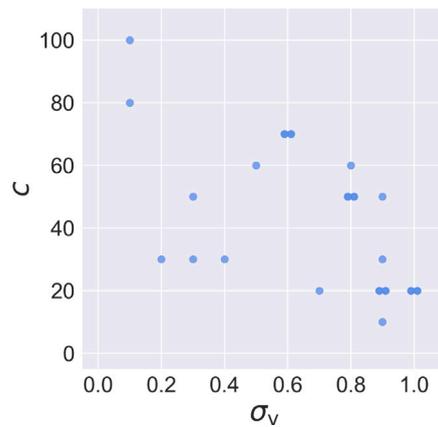


Fig. 9. Assessment of Fitted Parameters for the model with visual constraints and looming aversion assumptions (VLM). Points represent non-policy parameter combinations, with overlapping points displayed as slightly horizontally offset to indicate multiple occurrences.

Finally, we explored the influence of free parameters, σ_v and the looming aversion weight c , on the outcomes of our most effective model, the VLM. As illustrated in Fig. 8, we computed the mean CIT of VLM across all test scenarios for every combination of parameter values. There was a noticeable increase in CIT with a larger looming aversion weight c . We observed a reverse trend with regard to the visual noise parameter, σ_v . An increment in σ_v generally led to a decrease in CIT.

Fig. 9 presents a scatter plot that visualises two selected non-policy parameter combinations for the VLM model, with some data points slightly offset horizontally to denote multiple occurrences. The best-fit parameters lie inside the range of the search parameters, indicating that the selected parameter range we used was sufficiently broad.

3.3. Additional model tests

In addition to our main analyses, we conducted tests on four additional alternative models with different assumptions or different modelling approaches, but none of these performed better than the VLM model. Detailed results and figures from the tests of these alternative model variants can be found in Appendix A.2.

4. Discussion

4.1. Main findings

In this study, we developed a pedestrian crossing decision model. This model is based on the computational rationality framework and uses reinforcement learning (RL) to derive boundedly optimal behaviour policies under assumptions about human constraints and preferences. We see four main findings:

First, the VLM model variant, which incorporates both visual constraints and looming aversion assumptions, captured all four of our targeted behavioural phenomena. To our knowledge, no previous models have captured all four of these phenomena. In addition, these phenomena have not previously been addressed using a bounded optimality modelling framework.

Second, by testing different model variants, we demonstrated how specific human constraints affect pedestrian behaviour. In particular, our investigation of the model with visual constraints revealed that speed-dependent gap acceptance, previously interpreted as a ‘bias’ in decision-making, may instead reflect a rational adaptation to visual perception constraints.

Third, our research underscores the necessity of including both visual constraints and looming aversion assumptions. The model variants with only one of these assumptions, VM and LM, each only reproduced some of the phenomena. This finding highlights the complexity of pedestrian decision-making processes. Previous research, such as the study by Tian et al. (2022), has typically focused on single aspects of human constraints – like looming aversion – whereas our results suggest that both visual constraints and looming aversion play important roles, and both need to be integrated into the same model to explain the observed human behaviour.

Fourth, another important finding is the effectiveness of our parameter conditioning method. Applying this method to the VLM model allowed us to fit the model to the empirical data and demonstrate good quantitative fits with just two free model parameters per participant.

In the following subsections, we will discuss the impact of our mechanistic assumptions about human constraints and preferences on crossing decisions and explore our study’s broader implications and limitations.

4.2. Impact of mechanistic assumptions on crossing decisions

In this subsection, we will first discuss how assumptions about human constraints and preference affect the generation of the four targeted phenomena:

TTA-dependent gap acceptance. While both the model with the looming aversion assumption (LM) and the model with the visual constraints assumption (VM) demonstrate increased gap acceptance with rising initial TTA, this effect is more evident in the LM model compared to the VM model. This finding suggests that looming aversion plays an essential role in assessing safe crossing opportunities. In the LM model, a greater TTA reduces $\frac{1}{\tau}$, resulting in a lower looming aversion reward for crossing at higher TTA conditions. This phenomenon aligns with the model's inherent logic, where a higher TTA is associated with a lower perceived threat, encouraging the agent to cross.

Regarding VM, the qualitative pattern of TTA-dependent gap acceptance for this model is caused by the perceptual noise. Given that the agent cannot observe the vehicle's state accurately, there is a higher risk of collision during the crossing, especially in the lower TTA conditions. To avoid the penalty from the collision, the agent is more conservative and less likely to cross at the low initial TTA condition (2.3 s). However, at higher initial TTAs, the impact of the perceptual noise is less detrimental to making safe crossing decisions, leading to higher gap acceptance rates.

Speed-dependent gap acceptance. Our findings indicate that models incorporating the visual constraints assumption can generate speed-dependent gap acceptance behaviour. Understanding why speed-dependent gap acceptance arises from the visual constraints is not entirely trivial. Our best interpretation is that it is related to the more dispersed distribution of estimated TTA at lower vehicle speeds for a given time gap, as shown in Fig. A.14. Therefore, from a reward maximisation perspective, the agent makes more cautious decisions when interacting with a lower-speed vehicle.

This finding is interesting in relation to past explanations of speed-dependent gap acceptance as being caused by biases in human TTA estimation, due to humans making inappropriate use of distance information (Petzoldt, 2014) or visual looming information (Tian et al., 2022) when judging crossing safety. Our findings instead suggest that speed-dependent gap acceptance is an optimal behaviour, but that this is an optimality which is bounded by the particular noise characteristics of the human visual system. Differently put, given the nature of human visual perception, speed-dependent gap acceptance is an entirely rational behaviour. It is, of course, still possible that human pedestrians make use of visual looming information when making these decisions, as suggested by Tian et al. (2022), but if so, this seems to be part of a boundedly optimal strategy rather than a biased (i.e., suboptimal) heuristic.

Speed-dependent yielding acceptance. Our results show that both the model with the noisy perception assumption and the model with the looming aversion assumption can independently generate the speed-dependent yielding acceptance behaviour. This can be understood as follows: Under the noisy perception assumption, perceptual noise decreases as the distance between the vehicle and the agent reduces. When a vehicle approaches at a slower speed, for a given TTA, it is closer to the pedestrian, resulting in less perceptual noise. Consequently, the agent can more accurately estimate the vehicle's position, leading to an earlier decision to cross. For the model with looming aversion, the slower vehicle speed near the agent reduces the $\frac{1}{\tau}$ value. This lower value increases the reward for crossing, thus encouraging the agent to cross earlier during the vehicle's yielding.

Stopping distance-dependent yielding acceptance. This behaviour is captured effectively by the looming aversion assumption (but not by the noisy perception assumption). This can be understood by considering that when a vehicle decelerates from the same initial distance and speed, stopping at a greater deceleration, further away from the pedestrian, results in a smaller $\frac{1}{\tau}$, thereby reducing the penalty to the agent for crossing in front of the vehicle. Therefore, the agent's tendency to cross increases when the vehicle stops further away.

Exploration of non-policy parameters. As illustrated in Fig. 8, there is a noticeable increase in average *Crossing Initiation Time (CIT)* with a larger looming aversion weight c . This is reasonable since a higher value of c amplifies the perceived threat, thus leading the agent to make more cautious crossing decisions.

Conversely, higher noise levels σ_v in the visual system are associated with slight reductions in mean CIT. This outcome is somewhat counterintuitive, especially when contrasted with the smaller CIT exhibited by BM, which operates without visual noise. This disparity suggests a complex interaction between the looming aversion and visual noise parameters in VLM's decision-making process. A possible explanation could be that lower levels of visual noise allow the agent to predict future looming more accurately and, thus, more strategically delay crossing the road to reduce the looming aversion penalty during a yielding vehicle's approach. Conversely, when the level of visual noise is high, the agent's risk assessment of potential collisions becomes less reliable, which may lead to a tendency to cross the road more quickly to avoid the collision and time penalties despite the penalty for looming aversion, and thus maintain a lower overall penalty from the perspective of reward maximisation.

4.3. Implications and future work

In contrast to previous studies in modelling work (Pekkanen et al., 2021; Tian et al., 2022), we do not specify in our model how the agent should make decisions. Instead, our model allows the agent to learn boundedly optimal decision-making, given its environment and constraints. By continually encountering varied scenarios in the dynamic environment, the agent learns its policy through an iterative process of trial and error. It should be noted that this learning process does not in itself rely on any human data

(the VR data in this study was used only later, to select those non-policy parameter values which resulted in the most human-like RL policy).

This dynamic learning approach offers several advantages. For one, it imparts a level of adaptability to our model, enabling it to be attuned to different possible real-world conditions. Furthermore, given the great success of modern deep RL in learning behaviour policy for highly challenging tasks, the approach we have taken here should in principle be extensible to modelling road user behaviour and interaction also in more general, high-complexity traffic scenarios. This is an interesting direction for future work, which will require both the adoption of more advanced RL methods than we have used here, and importantly also further research to establish more complete mechanistic models of the involved human perceptual, cognitive, and motor constraints.

While our learning-based approach models the tendency of humans to naturally adapt and refine their judgements and actions based on their experiences, it is crucial to note that the RL algorithm by which our model learns boundedly optimal policy is not intended to mimic the learning process in humans. Our model converges towards a policy that could be similar to boundedly optimal human behaviour, but the path it takes to reach this convergence is almost certainly not human-like. The value of our RL approach lies in its ability to explore complex scenarios and generate behaviour that aligns with the principles of computational rationality, thereby providing a partly mechanistic and more explainable alternative to purely data-driven ML models.

Furthermore, our findings contribute to improving the predictive accuracy of pedestrian crossing models. We reveal that two parameters – the looming aversion weight and the noise of agents' visual systems – have an important influence on the modelled CIT. In this regard, using real-world data from a large representative population to ascertain these parameters and subsequently incorporating them into RL processes may be important steps towards enhanced model accuracy.

The method of training the RL model conditioned on non-policy parameters is established in computational rationality within the human-computer interaction domain (Keurulainen et al., 2023). Here, we demonstrated that we could successfully apply this method to pedestrian modelling. This parameter-conditioning is particularly beneficial if the parameter space is large, so an interesting future direction will be to investigate pedestrian models conditioned also on a wider range of non-policy parameters. Again, applying this model to real-world pedestrian data presents a promising research direction. For example, we could infer non-policy parameters, such as risk perception or walking speed, which are unique to different individuals, which could allow more precise prediction of pedestrian behaviours in various traffic scenarios.

As mentioned in Section 2.1, while the sample size of only 20 participants might seem small, it is important to note that our main findings concern the model's ability to replicate the observed qualitative phenomena, and that these have been observed in many other datasets, including naturalistic datasets (Oxley et al., 2005; Lobjois and Cavallo, 2007; Petzoldt, 2014; Tian et al., 2022; Ackermann et al., 2019; Tian et al., 2023). This alignment across different studies supports the validity of our findings despite the limited sample. Furthermore, the main direct use of the dataset here is to demonstrate that we can do parameter inference with our approach, including fitting to individual participants. It is important to note that we are not suggesting that the specific parameter values we have obtained here will necessarily generalise, e.g., to larger naturalistic samples. Instead, we reproduce qualitative behaviour phenomena that have also been observed in naturalistic data, and we show that our model can be fitted quantitatively to data. In future work, we plan to extend our model to address naturalistic datasets where more diverse scenarios are available. This can not only improve the applicability of our model to real-world autonomous vehicle (AV) applications but also increase the generalisability of the observed behaviour.

In this study, we used VR to simulate real-world conditions. VR offers a controlled, safe, and replicable environment that is helpful for studying pedestrian behaviour. Kalantari et al. (2023) have shown that pedestrian crossing behaviours observed in VR settings are comparable to those in real-world environments, lending credence to the use of VR as a reliable method for behavioural studies. We believe this is an important step towards applying this modelling method to more complex real-world scenarios.

The scenario involving one pedestrian crossing at an unsignalised road was specifically chosen to isolate and model pedestrian reactions to one oncoming vehicle without the influence of traffic control devices such as traffic lights. This specific focus was selected to simplify the crossing environment, ensuring that the observed behaviours were predominantly influenced by the vehicle's approach to provide a clear view of the factors influencing pedestrian behaviour under these conditions. In future work, we could include more complex scenarios, such as roads with traffic lights or interactions between pedestrians and multiple vehicles, to enhance the practical applicability of our research.

In addition, the scenarios presented here focused primarily on constant-speed and deceleration conditions. Although these two scenarios are commonly seen in pedestrian-vehicle interaction, other scenarios may also be interesting; for example, drivers may accelerate when refusing to give pedestrians priority (Markkula et al., 2023). To include more diverse scenarios in future studies will be helpful for a more comprehensive pedestrian behaviour model.

In the field of AV simulation, the potential applications of our model are promising. For instance, it could improve the realism and effectiveness of AV testing simulations (Markkula and Dogar, 2022). Furthermore, by incorporating a detailed understanding of pedestrian behaviour, real-time AV algorithms can be fine-tuned to anticipate and respond to pedestrian actions more accurately. For example, by combining our model's predictions with real-time data from vehicle sensors, AV systems could achieve a higher level of situational awareness, enabling them to make safer decisions in complex urban environments (Camara et al., 2020). It should be noted that it utilises only pedestrian and vehicle trajectory information—specifically, their positions and velocities. This type of data is readily captured by the sensor suites in modern vehicles for real-time sensing. Similar uses of these models in, for example, infrastructure design can also benefit pedestrian traffic safety more broadly, also in conventional human-driven traffic (Kraidi and Evdorides, 2020; Zhu et al., 2022).

Regarding the real-world applicability of our model in autonomous vehicle contexts, it is important to note that it utilises only pedestrian and vehicle trajectory information—specifically, their positions and velocities. This type of data is readily captured by the sensor suites in modern vehicles for real-time sensing.

There is clear and ample scope for further improvements to our model. For example, there is more variability in the human crossing initiation time than in the model. This could be due to both between- and within-individual variability of decision-influencing factors. In addition, our reward function is simple, and so far, we have fine-tuned only parts of it (the looming aversion). This was sufficient for our purposes, to show qualitative patterns of human road-crossing, and contrary to our expectations, it even allowed us to achieve relatively good quantitative fits, but there is clear room for further work. For example, following a similar approach to what we did here with σ_v , we could also tune, for example, the time-loss penalty in the reward function. This adjustment would allow the model to more accurately reflect how different pedestrians value the trade-off between time and safety.

Another aspect to consider for future improvement of our model is its current focus exclusively on the pedestrian's reaction to vehicles, without accounting for the driver's reaction to the pedestrian as discussed in Section 2. Real-world scenarios often involve a dynamic interplay between both parties, influencing each other's decisions. To address this gap, future studies could employ multiagent RL to model the interactions between pedestrians and drivers more realistically (Hu et al., 1998).

Finally, it is interesting to note that we face the challenge of fully understanding the complexity of our model's behaviour. This issue arises both in relation to the speed-dependent gap acceptance (we have shown that this behaviour is boundedly optimal, but we have not been able to fully explain why) and the effect of visual noise levels on CITs (zero noise causes rapid decisions but low levels of noise cause slower decisions than high levels of noise). While our computational rationality model is more interpretable than fully data-driven ML models, it only partially eliminates the challenge of explicating the 'why' behind its decision-making processes. On a high level, the behaviour of our model, governed by the principle of computational rationality and informed by RL, can be explained quite simply as boundedly optimal, but formulating a more detailed explanation of *why* a certain behaviour is boundedly optimal may not always be straightforward.

5. Conclusions

In conclusion, this study contributes to modelling human pedestrian crossing decisions, utilising a computational rationality framework, and using RL to identify boundedly optimal behaviour policy. Under our chosen human-like constraints, we demonstrate that our model can emulate human-like behaviour in road-crossing scenarios with both non-yielding and yielding vehicles. Interestingly, our findings suggest that the previously reported speed-dependent gap acceptance behaviour in pedestrian decision-making can be understood as a rational adaptation to visual perception constraints. Our best model not only reproduces this behaviour but also qualitatively and quantitatively captures three other key phenomena: TTA-dependent gap acceptance, speed-dependent yielding acceptance, and stopping distance-dependent yielding acceptance. In addition, we identified two parameters – the weight of looming aversion and the noise in agents' visual systems – which can provide a better understanding of human decision-making in real-world pedestrian crossing scenarios. While we recognise the need to investigate and develop our model further, the insight from this study provides a promising foundation for improved pedestrian models. Not least, the inherent adaptability of RL suggests the potential for future extensions to more complex traffic scenarios, a challenge that traditional mechanistic models may find difficult to address. The potential applications of our model are notable. In the field of AVs, our results can contribute to more realistic simulated road users in virtual testing environments and guide vehicular decision-making algorithms to better anticipate human behaviour, ensuring safer human-AV co-existence.

CRedit authorship contribution statement

Yueyang Wang: Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Aravinda Ramakrishnan Srinivasan:** Writing – review & editing, Validation, Supervision, Methodology, Investigation, Formal analysis, Conceptualization. **Jussi P.P. Jokinen:** Writing – review & editing, Methodology, Conceptualization. **Antti Oulasvirta:** Writing – review & editing, Methodology, Conceptualization. **Gustav Markkula:** Writing – review & editing, Validation, Supervision, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

This work was supported by the UK Engineering and Physical Sciences Research Council (grant EP/S005056/1); the Academy of Finland (grant 330347); and the University of Leeds International Academic Mobility Scheme. We thank the four anonymous reviewers, whose comments have greatly improved the quality of our paper.

Appendix. Additional material

A.1. Model convergence

Fig. A.10 illustrates the average rewards of four main model variants recorded for each of the 500 episodes, spanning the entirety of the training phase. From Fig. A.10, it can be noted that the reward for the models, including the looming aversion assumption (LM and VLM), achieve lower rewards overall due to the change in the reward structure introduced by this assumption. Additionally, adding visual constraints to the baseline model (going from BM to VM) led to a decline in the average rewards. This is due to the VM agent's inability to know the precise position of the approaching vehicle. As a result, the agent tends to make more cautious decisions, leading to increased time penalties. However, there was no noticeable difference in the achievable rewards for LM and VLM. A possible reason could be that for LM, the looming aversion penalty is directly proportional to the inverse τ , while in VLM, the looming aversion penalty corresponds to the estimated inverse τ . From the reward maximisation perspective, both model variants aim to minimise this looming penalty within their respective environmental contexts, potentially explaining the similar overall rewards.

A.2. Additional model variants

Beyond the core model variants outlined in Section 2.2, we explored four additional variations to deepen our understanding of the pedestrian decision-making process:

- **VLDM:** The model variant considers motor delay an additional human-like constraint, reflecting the time lag between decision-making and action execution. This variant expands upon the VLM model variant, incorporating motor delay into the model.
- **VNM:** A variation of the looming aversion model that tests an alternative representation of risk aversion. Instead of an inverse τ , we evaluated a model where aversion is formulated against near-collision situations, providing an alternative perspective on risk assessment in pedestrian behaviour.
- **VLM (E):** A version of the VLM model that employs a single set of non-policy parameters across the entire human dataset instead of fitting them to each participant. This approach aimed to assess the efficacy of a more generalised model application.
- **VLM (S):** This variation of the VLM model involves training individual networks for each unique set of parameters.

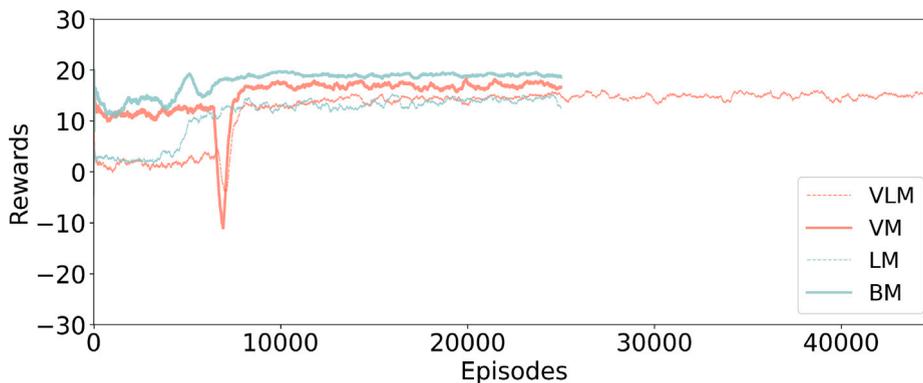


Fig. A.10. Convergence plot for different models.

Table A.5
Quantitative assessment of additional model performance.

Model type	Log Lik.	Params.	AIC	MAD (s)
VLDM	-536	60	1192	0.35
VNM	-564	40	1208	0.99
VLM(E)	-547	2	1098	0.36
VLM(S)	-572	40	1225	0.39

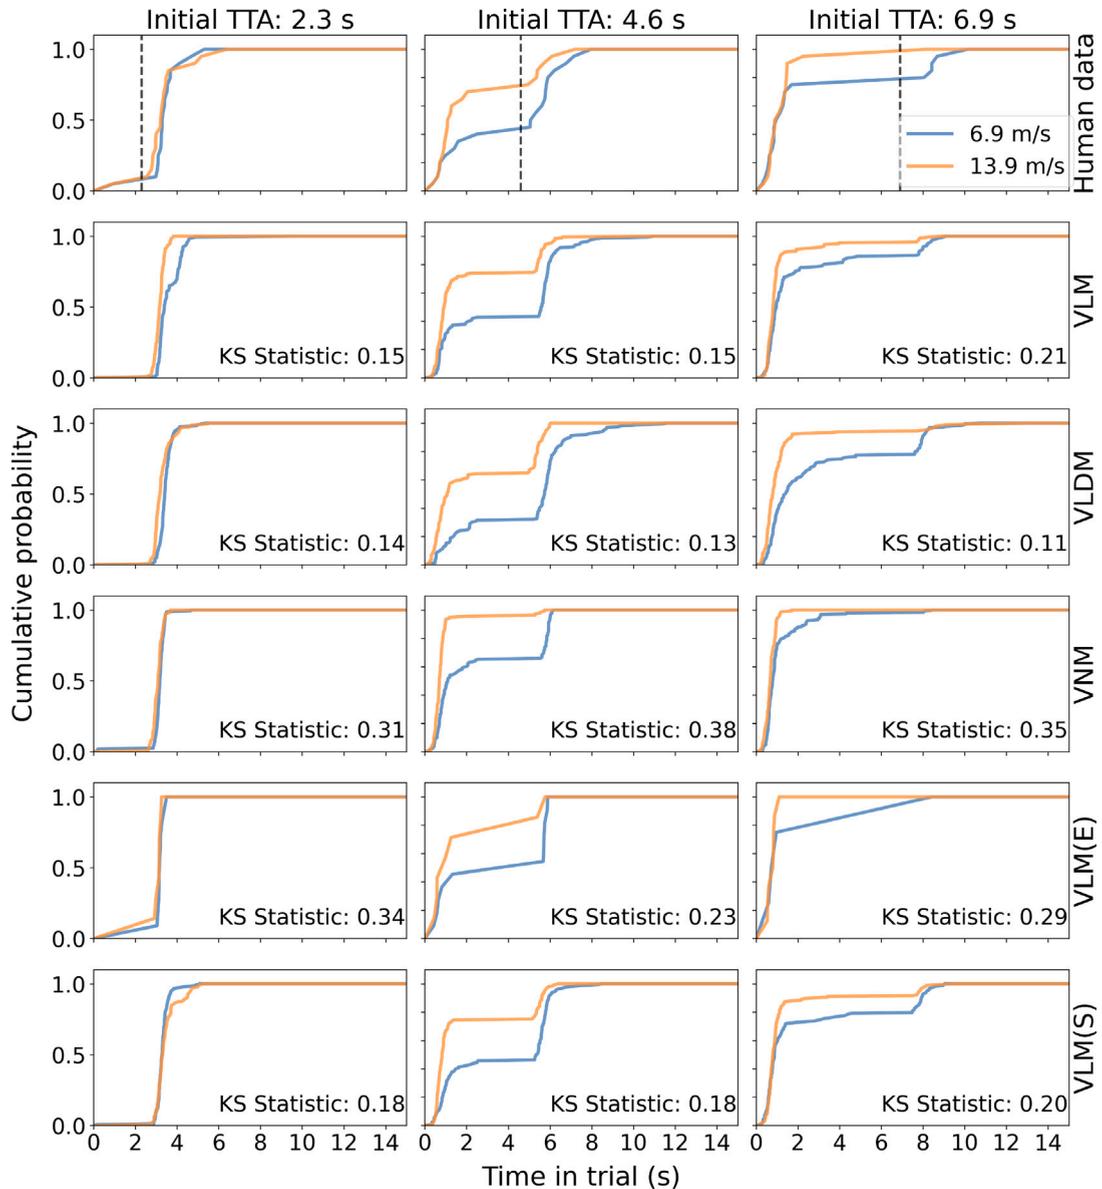


Fig. A.11. Cumulative probability for Crossing Initiation Time (CIT) in constant speed scenarios for human data and additional models. Black dashed vertical lines in Human data indicate the times vehicles passed pedestrians. To avoid repetition, these black lines are not included in the model results. Vehicles passed pedestrians at the same time in the trial for each initial TTA condition. The VLM model is included in this figure for comparison.

These models were designed to assess the influence of additional factors and assumptions on pedestrian crossing behaviour. While these models provided valuable insights, they did not show an overall improvement in performance compared to our primary model, VLM. Figs. A.11–A.13 and Table A.5 show the results obtained for these models.

The VLDM model integrates motor delay (m) as an additional non-policy parameter, alongside visual constraints (σ_v) and looming aversion (c). VLDM showed similar performance to the VLM model regarding log-likelihood and MAD. However, it obtained a higher AIC score due to the increased model complexity.

The VNM model explores the concept of pedestrian aversion to near-collision situations. This model variant provides an alternative way of modelling pedestrian aversion to being in safety-critical situations as an alternative to the formulation based on looming in the LM/VLM models. This model assumes pedestrians inherently aim to avoid situations where a vehicle invades

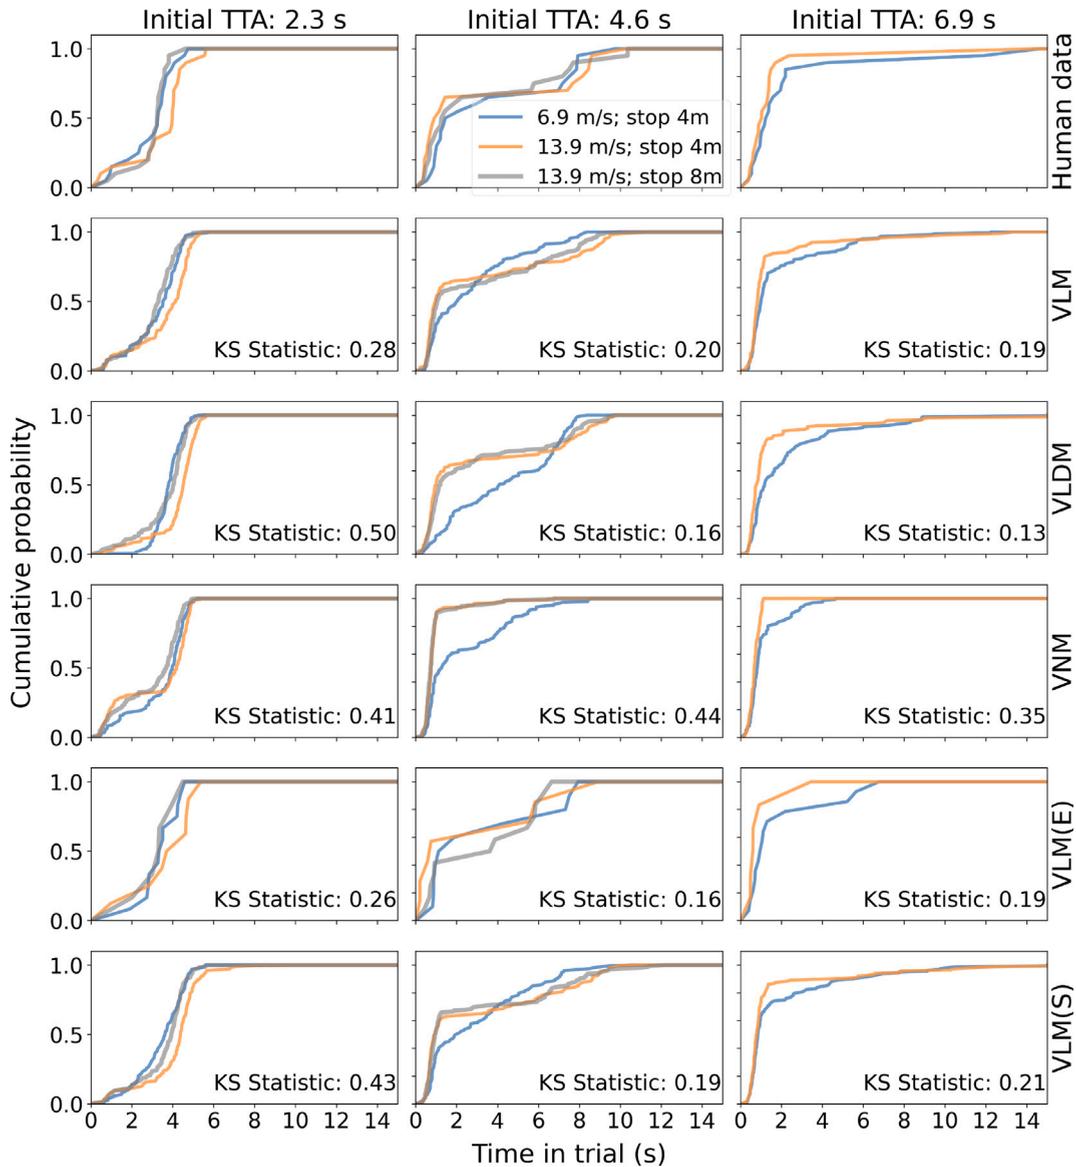


Fig. A.12. Cumulative probability for CIT in yielding scenarios for human data and additional models. The VLM model is included in this figure for comparison.

their personal space, leading to an additional penalty in the model when this occurs. While VNM exhibited more human-like CITs than the VM model, it still lagged behind the comprehensive performance of the VLM model regarding AIC. This disparity in model performance could imply that the integration of temporal-spatial clues (e.g., looming) may play a more important role in pedestrian decision-making than spatial factors alone (e.g., proximity to a near-collision zone).

The VLM (E) model, which uses a single set of non-policy parameters for the entire dataset, performed better in terms of AIC value than VLM, reflecting a simpler model structure. However, it did not match the log-likelihood performance of VLM. This finding suggests that individualising non-policy parameters may be beneficial for accurately capturing individual pedestrian behaviour.

In the VLM (S) model, instead of conditioning a single RL policy on the values of the non-policy parameters, we trained separate RL policy networks for each of the $10 \times 10 = 100$ considered values of the non-policy parameters. Despite the increased computational effort, VLM (S) did not achieve better performance metrics than the original VLM model, as shown in Table A.5. This result highlights the efficiency of our chosen parameter-conditioning in the VLM model.

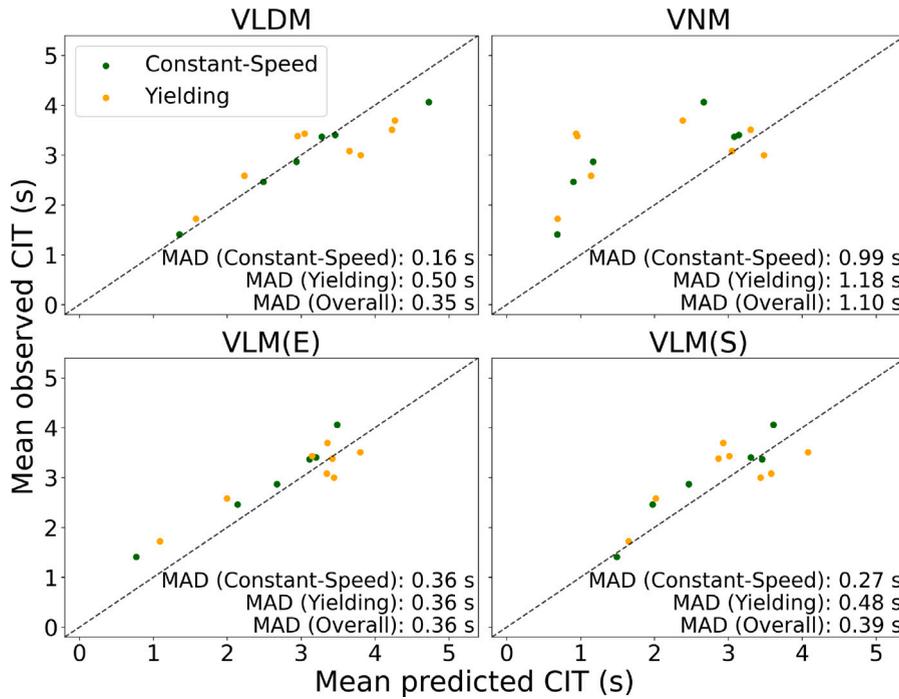


Fig. A.13. Predicted vs. observed mean CIT across the different scenarios.

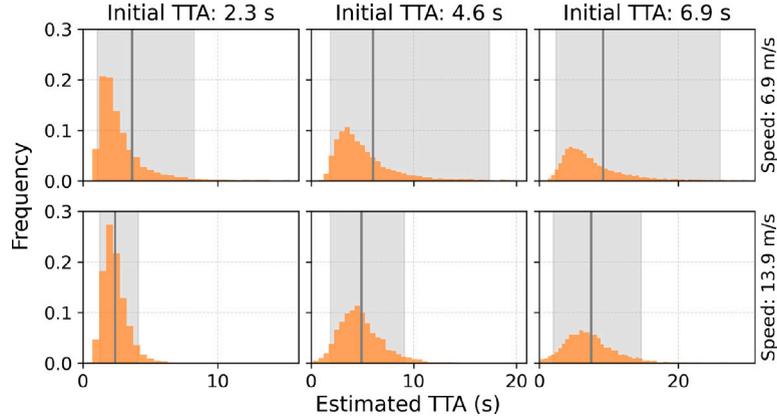


Fig. A.14. Distributions of initial estimated TTA from the output of the visual perception model, across the different scenarios in the experiment. The grey vertical line is the mean value, and the shading area shows the 5th and 95th percentiles of the estimated TTA.

Regarding VLDM, with the motor delay, m , as an additional input parameter alongside σ_v and c , from Table A.5, we can find that this model matched the VLM's performance in terms of both log-likelihood and MAD. However, its additional parameters result in a higher AIC compared to VLM.

Data availability

All data, and models and code are available at <https://github.com/YYWang98/Pedestrian-crossing-decisions.git>.

References

- Abughalieh, K.M., Alawneh, S.G., 2020. Predicting pedestrian intention to cross the road. *IEEE Access* 8, 72558–72569.
- Ackermann, C., Beggiato, M., Bluhm, L.-F., Löw, A., Krems, J.F., 2019. Deceleration parameters and their applicability as informal communication signal between pedestrians and automated vehicles. *Transp. Res. F* 62, 757–768.
- Al-Shihabi, T., Mourant, R.R., 2001. A framework for modeling human-like driving behaviors for autonomous vehicles in driving simulators. In: *Proceedings of the Fifth International Conference on Autonomous Agents*. pp. 286–291.
- Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S., 2016. Social lstm: Human trajectory prediction in crowded spaces. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 961–971.
- Althoff, M., Lutz, S., 2018. Automatic generation of safety-critical test scenarios for collision avoidance of road vehicles. In: *2018 IEEE Intelligent Vehicles Symposium. IV, IEEE*, pp. 1326–1333.
- Asaithambi, G., Kuttan, M.O., Chandra, S., 2016. Pedestrian road crossing behavior under mixed traffic conditions: a comparative study of an intersection before and after implementing control measures. *Transp. Dev. Econ.* 2 (2), 14.
- Banovic, N., Oulasvirta, A., Kristensson, P.O., 2019. Computational modeling in human-computer interaction. In: *Extended Abstracts of the 2019 CHI Conference on Human Factors in Computing Systems*. pp. 1–7.
- Camara, F., Bellotto, N., Cosar, S., Weber, F., Nathanael, D., Althoff, M., Wu, J., Ruenz, J., Dietrich, A., Markkula, G., et al., 2020. Pedestrian models for autonomous driving part ii: high-level models of human behavior. *IEEE Trans. Intell. Transp. Syst.* 22 (9), 5453–5472.
- Chen, X., Acharya, A., Oulasvirta, A., Howes, A., 2021b. An adaptive model of gaze-based selection. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. pp. 1–11.
- Chen, H., Chang, H.J., Howes, A., 2021a. Apparently irrational choice as optimal sequential decision making. In: *Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 35*, pp. 792–800.
- Crosato, L., Tian, K., Shum, H.P., Ho, E.S., Wang, Y., Wei, C., 2023. Social interaction-aware dynamical models and decision-making for autonomous vehicles. *Adv. Intell. Syst.* 2300575.
- Dai, S., Li, L., Li, Z., 2019. Modeling vehicle interactions via modified LSTM models for trajectory prediction. *IEEE Access* 7, 38287–38296.
- DeLucia, P.R., 2008. Critical roles for distance, task, and motion in space perception: Initial conceptual framework and practical implications. *Hum. Factors* 50 (5), 811–820.
- DeLucia, P.R., 2015. Perception of collision. In: Hoffman, R.R., Hancock, P.A., Scerbo, M.W., Parasuraman, R., Szalma, J.L. (Eds.), *The Cambridge Handbook of Applied Perception Research*. In: *Cambridge Handbooks in Psychology*, Cambridge University Press, pp. 568–591. <http://dx.doi.org/10.1017/CBO9780511973017.035>.
- Diaz-Ruiz, C.A., Xia, Y., You, Y., Nino, J., Chen, J., Monica, J., Chen, X., Luo, K., Wang, Y., Emond, M., et al., 2022. Ithaca365: Dataset and driving perception under repeated and challenging weather conditions. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 21383–21392.
- Doellinger, J., Spies, M., Burgard, W., 2018. Predicting occupancy distributions of walking humans with convolutional neural networks. *IEEE Robot. Autom. Lett.* 3 (3), 1522–1528.
- Domeyer, J., Dinparastdjadid, A., Lee, J.D., Douglas, G., Alsaïd, A., Price, M., 2019. Proxemics and kinesics in automated vehicle–pedestrian communication: Representing ethnographic observations. *Transp. Res. Rec.* 2673 (10), 70–81.
- Duric, Z., Gray, W.D., Heishman, R., Li, F., Rosenfeld, A., Schoelles, M.J., Schunn, C., Wechsler, H., 2002. Integrating perceptual and cognitive modeling for adaptive and intelligent human-computer interaction. *Proc. IEEE* 90 (7), 1272–1289.
- Faisal, A.A., Selen, L.P., Wolpert, D.M., 2008. Noise in the nervous system. *Nat. Rev. Neurosci.* 9 (4), 292–303.
- Fajen, B.R., Warren, W.H., 2003. Behavioral dynamics of steering, obstacle avoidance, and route selection. *J. Exp. Psychol. [Hum. Percept.]* 29 (2), 343.
- Gershman, S.J., Horvitz, E.J., Tenenbaum, J.B., 2015. Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science* 349 (6245), 273–278.
- Giles, O., Markkula, G., Pekkanen, J., Yokota, N., Matsunaga, N., Merat, N., Daimon, T., 2019. At the zebra crossing: Modelling complex decision processes with variable-drift diffusion models. In: *Proceedings of the 41st Annual Meeting of the Cognitive Science Society*. Cognitive Science Society, pp. 366–372.
- González-Méndez, M., Olaya, C., Fasolino, I., Grimaldi, M., Obregón, N., 2021. Agent-based modeling for urban development planning based on human needs. Conceptual basis and model formulation. *Land Use Policy* 101, 105110.
- Gorrini, A., Crociani, L., Vizzari, G., Bandini, S., 2018. Observation results on pedestrian-vehicle interactions at non-signalized intersections towards simulation. *Transp. Res. F* 59, 269–285.
- Howes, A., Jokinen, J.P., Oulasvirta, A., 2023. Towards machines that understand people. *AI Mag.* 44 (3), 312–327.
- Hu, J., Wellman, M.P., et al., 1998. Multiagent reinforcement learning: theoretical framework and an algorithm. In: *ICML. Vol. 98*, pp. 242–250.
- Jain, A., Gupta, A., Rastogi, R., 2014. Pedestrian crossing behaviour analysis at intersections. *Int. J. Traffic Transp. Eng.* 4 (1), 103–116.
- Jokinen, J.P., Kujala, T., Oulasvirta, A., 2021. Multitasking in driving as optimal adaptation under uncertainty. *Hum. Factors* 63 (8), 1324–1341.
- Jokinen, J., Remes, U., Kujala, T., Corander, J., 2022. Bayesian parameter inference for cognitive simulators. In: Williamson, J.H., Oulasvirta, A., Kristensson, P.O., Banovic, N. (Eds.), *Bayesian Methods for Interaction and Design*. Cambridge University Press, Cambridge, United Kingdom, pp. 308–334. <http://dx.doi.org/10.1017/9781108874830.016>.
- Jokinen, J.P., Wang, Z., Sarcar, S., Oulasvirta, A., Ren, X., 2020. Adaptive feature guidance: Modelling visual search with graphical layouts. *Int. J. Hum.-Comput. Stud.* 136, 102376.
- Kalantari, A.H., Lin, Y., Mohammadi, A., Merat, N., Markkula, G., 2023. Testing the validity of multi participant distributed simulation for understanding and modeling road user interaction. <http://dx.doi.org/10.31234/osf.io/gk9af>, URL: osf.io/preprints/psyarxiv/gk9af.
- Keurulainen, A., Westerlund, I.R., Keurulainen, O., Howes, A., 2023. Amortised experimental design and parameter estimation for user models of pointing. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. pp. 1–17.
- Klischat, M., Liu, E.L., Holtke, F., Althoff, M., 2020. Scenario factory: Creating safety-critical traffic scenarios for automated vehicles. In: *2020 IEEE 23rd International Conference on Intelligent Transportation Systems. ITSC, IEEE*, pp. 1–7.
- Knill, D.C., Pouget, A., 2004. The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends Neurosci.* 27 (12), 712–719.
- Kraïdi, R., Evdorides, H., 2020. Pedestrian safety models for urban environments with high roadside activities. *Saf. Sci.* 130, 104847.
- Kumamoto, K., Yamada, K., 2017. CNN-based pedestrian orientation estimation from a single image. In: *2017 4th IAPR Asian Conference on Pattern Recognition. ACPR, IEEE*, pp. 13–18.
- Kwon, O.-S., Tadin, D., Knill, D.C., 2015. Unifying account of visual motion and position perception. *Proc. Natl. Acad. Sci.* 112 (26), 8142–8147.
- Leu, G., Curtis, N.J., Abbass, H., 2012. Modeling and evolving human behaviors and emotions in road traffic networks. *Procedia-Soc. Behav. Sci.* 54, 999–1009.
- Lewis, R.L., Howes, A., Singh, S., 2014. Computational rationality: Linking mechanism and behavior through bounded utility maximization. *Top. Cogn. Sci.* 6 (2), 279–311.
- Li, Z., Ko, Y.-J., Putkonen, A., Feiz, S., Ashok, V., Ramakrishnan, I., Oulasvirta, A., Bi, X., 2023. Modeling touch-based menu selection performance of blind users via reinforcement learning. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. CHI '23*, Association for Computing Machinery, New York, NY, USA, p. 18. <http://dx.doi.org/10.1145/3544548.3580640>.
- Li, L., Liu, Y., Wang, J., Deng, W., Oh, H., 2016. Human dynamics based driver model for autonomous car. *IET Intell. Transp. Syst.* 10 (8), 545–554.

- Li, L., Pagnucco, M., Song, Y., 2022. Graph-based spatial transformer with memory replay for multi-future pedestrian trajectory prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2231–2241.
- Lieder, F., Griffiths, T.L., 2020. Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behav. Brain Sci.* 43, e1.
- Lobjois, R., Cavallo, V., 2007. Age-related differences in street-crossing decisions: The effects of vehicle speed and time constraints on gap selection in an estimation task. *Accid. Anal. Prev.* 39 (5), 934–943.
- Lorenzo, J., Alonso, I.P., Izquierdo, R., Ballardini, A.L., Saz, Á.H., Llorca, D.F., Sotelo, M.Á., 2021. Capformer: Pedestrian crossing action prediction using transformer. *Sensors* 21 (17), 5694.
- Markkula, G., Dogar, M.R., 2022. Models of human behavior for human-robot interaction and automated driving: How accurate do the models of human behavior need to be? *IEEE Robotics I Autom. Mag.* 2–7. <http://dx.doi.org/10.1109/MRA.2022.3182892>.
- Markkula, G., Engström, J., Lodin, J., Bårgman, J., Victor, T., 2016. A farewell to brake reaction times? Kinematics-dependent brake response in naturalistic rear-end emergencies. *Accid. Anal. Prev.* 95, 209–226.
- Markkula, G., Lin, Y.-S., Srinivasan, A.R., Billington, J., Leonetti, M., Kalantari, A.H., Yang, Y., Lee, Y.M., Madigan, R., Merat, N., 2023. Explaining human interactions on the road by large-scale integration of computational psychological theory. *PNAS Nexus* 2 (6), 163.
- Ooi, T.L., Wu, B., He, Z.J., 2001. Distance determined by the angular declination below the horizon. *Nature* 414 (6860), 197–200.
- Oulasvirta, A., Jokinen, J.P., Howes, A., 2022. Computational rationality as a theory of interaction. In: Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems. pp. 1–14.
- Oxley, J.A., Ihsen, E., Fildes, B.N., Charlton, J.L., Day, R.H., 2005. Crossing roads safely: an experimental study of age differences in gap selection by pedestrians. *Accid. Anal. Prev.* 37 (5), 962–971.
- Pekkanen, J., Giles, O.T., Lee, Y.M., Madigan, R., Daimon, T., Merat, N., Markkula, G., 2021. Variable-drift diffusion models of pedestrian road-crossing decisions. *Comput. Brain Behav.* 1–21.
- Pelorusso, R., 2020. Modeling and urban planning: A systematic review of performance-based approaches. *Sustain. Cities Soc.* 52, 101867.
- Petzoldt, T., 2014. On the relationship between pedestrian gap acceptance and time to arrival estimates. *Accid. Anal. Prev.* 72, 127–133.
- Quan, R., Zhu, L., Wu, Y., Yang, Y., 2021. Holistic LSTM for pedestrian trajectory prediction. *IEEE Trans. Image Process.* 30, 3229–3239.
- Ratcliff, R., Smith, P.L., Brown, S.D., McKoon, G., 2016. Diffusion decision model: Current issues and history. *Trends Cogn. Sci.* 20 (4), 260–281.
- Risto, M., Emmenegger, C., Vinkhuyzen, E., Cefkin, M., Hollan, J., 2017. Human-vehicle interfaces: The power of vehicle movement gestures in human road user coordination. In: Driving Assessment Conference. Vol. 9, University of Iowa.
- Schneemann, F., Gohl, I., 2016. Analyzing driver-pedestrian interaction at crosswalks: A contribution to autonomous driving in urban environments. In: 2016 IEEE Intelligent Vehicles Symposium. IV, IEEE, pp. 38–43.
- Silver, D., Huang, A., Maddison, C.J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al., 2016. Mastering the game of go with deep neural networks and tree search. *Nature* 529 (7587), 484–489.
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., Lai, M., Bolton, A., et al., 2017. Mastering the game of go without human knowledge. *Nature* 550 (7676), 354–359.
- Simon, H.A., 1955. A behavioral model of rational choice. *Q. J. Econ.* 99–118.
- Srinivasan, A.R., Lin, Y.-S., Antonello, M., Knittel, A., Hasan, M., Hawasly, M., Redford, J., Ramamoorthy, S., Leonetti, M., Billington, J., et al., 2023. Beyond RMSE: Do machine-learned models of road user interaction produce human-like behavior? *IEEE Trans. Intell. Transp. Syst.*
- Stocker, A.A., Simoncelli, E.P., 2006. Noise characteristics and prior expectations in human visual speed perception. *Nature Neurosci.* 9 (4), 578–585.
- Sun, R., Zhuang, X., Wu, C., Zhao, G., Zhang, K., 2015. The estimation of vehicle speed and stopping distance by pedestrians crossing streets in a naturalistic traffic environment. *Transp. Res. F* 30, 97–106.
- Sutton, R.S., Barto, A.G., 2018. Reinforcement Learning: An Introduction. MIT Press.
- Tian, K., Markkula, G., Wei, C., Lee, Y.M., Madigan, R., Merat, N., Romano, R., 2022. Explaining unsafe pedestrian road crossing behaviours using a psychophysics-based gap acceptance model. *Saf. Sci.* 154, 105837.
- Tian, K., Tzigieras, A., Wei, C., Lee, Y.M., Holmes, C., Leonetti, M., Merat, N., Romano, R., Markkula, G., 2023. Deceleration parameters as implicit communication signals for pedestrians' crossing decisions and estimations of automated vehicle behaviour. *Accid. Anal. Prev.* 190, 107173.
- Turnwald, A., Althoff, D., Wollherr, D., Buss, M., 2016. Understanding human avoidance behavior: interaction-aware decision making based on game theory. *Int. J. Soc. Robotics* 8, 331–351.
- Van Hasselt, H., Guez, A., Silver, D., 2016. Deep reinforcement learning with double q-learning. In: Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 30.
- Von Neumann, J., Morgenstern, O., 1947. Theory of games and economic behavior, 2nd rev.
- Wang, Y., Ramakrishnan Srinivasan, A., Jokinen, J.P., Oulasvirta, A., Markkula, G., 2023. Modeling human road crossing decisions as reward maximization with visual perception limitations. In: 2023 IEEE Intelligent Vehicles Symposium. IV, pp. 1–6. <http://dx.doi.org/10.1109/IV55152.2023.10186617>.
- Wang, Z., Schaul, T., Hessel, M., Hasselt, H., Lanctot, M., Freitas, N., 2016. Dueling network architectures for deep reinforcement learning. In: International Conference on Machine Learning. PMLR, pp. 1995–2003.
- Wunder, M., Littman, M.L., Babes, M., 2010. Classes of multiagent q-learning dynamics with epsilon-greedy exploration. In: Proceedings of the 27th International Conference on Machine Learning (ICML-10). pp. 1167–1174.
- Xie, Z., Clary, P., Dao, J., Morais, P., Hurst, J., van de Panne, M., 2019. Iterative reinforcement learning based design of dynamic locomotion skills for cassie. *arXiv preprint arXiv:1903.09537*.
- Xu, Y., Yang, X., Gong, L., Lin, H.-C., Wu, T.-Y., Li, Y., Vasconcelos, N., 2020. Explainable object-induced action decision for autonomous vehicles. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9523–9532.
- Yang, D., Özgüner, Ü., Redmill, K., 2020. A social force based pedestrian motion model considering multi-pedestrian interaction with a vehicle. *ACM Trans. Spatial Algorithms Syst. (TSAS)* 6 (2), 1–27.
- Yi, S., Li, H., Wang, X., 2016. Pedestrian behavior understanding and prediction with deep neural networks. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, the Netherlands, October 11–14, 2016, Proceedings, Part I 14. Springer, pp. 263–279.
- Yin, Z., Liu, R., Xiong, Z., Yuan, Z., 2021. Multimodal transformer networks for pedestrian trajectory prediction. In: IJCAI. pp. 1259–1265.
- Yuan, Y., Weng, X., Ou, Y., Kitani, K.M., 2021. Agentformer: Agent-aware transformers for socio-temporal multi-agent forecasting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9813–9823.
- Zhang, X., Angeloudis, P., Demiris, Y., 2022. ST CrossingPose: A spatial-temporal graph convolutional network for skeleton-based pedestrian crossing intention prediction. *IEEE Trans. Intell. Transp. Syst.* 23 (11), 20773–20782.
- Zhu, M., Sze, N., Newnam, S., 2022. Effect of urban street trees on pedestrian safety: A micro-level pedestrian casualty model using multivariate Bayesian spatial approach. *Accid. Anal. Prev.* 176, 106818.