Cognitive Mechanisms of Behaviour Estimation: Modelling Pedestrian Interpretation of Approaching Vehicle Behaviour

Athanasios Tzigieras

Submitted in accordance with the requirements for the degree of

Doctor of Philosophy

The University of Leeds

Institute for Transport Studies

June 2025

Intellectual Property and Publications

The candidate confirms that the work submitted is his own, except where work which has formed part of a jointly authored publication has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

One publication has been produced from research that was undertaken as part of this thesis. This publication is listed below with its full reference and details of the locations within this thesis.

Part of the work in Chapter 2 of this thesis, which details the first experiment and the collected dataset, has appeared in the following publication. This dataset was also subsequently used for the analyses in Chapter 3.

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. *Accident Analysis & Prevention*, 190, 107173.

The candidate's contributions to this publication included the conceptualisation, methodology, data curation, and writing – review & editing. Under the guidance and review of GM, NM, ML, CW, RR and CH, the candidate was responsible for the initial conceptualisation of the study and the development of the experimental methodology alongside the first author, KT. The candidate also contributed to the data collection, the data curation process and the review and editing of the final manuscript. The formal analysis and the writing of the original draft were led by the first author, KT. The manuscript was improved by comments from all the coauthors. Note that both Chapter 2 and Chapter 3 of this thesis use the collected data for different purposes than in the journal paper.

Acknowledgements

To my supervisors, Gustav Markkula, Natasha Merat, and Matteo Leonetti, I must express my deepest appreciation. Your invaluable guidance, support, patience, and understanding, even in the most challenging moments, meant everything. Gustav, I could not have embarked on, let alone completed, this PhD journey without your mentorship.

I would like to extend my gratitude to my examiners, Dr Thomas Hancock, Professor Martin Baumann and Professor Samantha Jamson. I am thankful for the stimulating and insightful discussions with them during the examination process.

This research was made possible by the generous funding from Nissan. I extend my sincere and heartfelt gratitude to Bob Bateman and Christopher Holmes for entrusting me with this wonderful opportunity.

A special thank you is owed to my friend and research partner, Kai Tian, for his strong work ethic and for the great times and insightful discussions we had. I am very fortunate to have worked with you.

My ideas would have remained just that without the amazing work of the simulator team, and so, I am grateful to Jorge Garcia de Pedro, whose excellent work in HIKER brought my experiments to life.

I feel incredibly lucky to have been surrounded by such excellent friends and colleagues in the Human Factors and Safety research group – Yueyang Wang, Chen Peng, Yee Mun Lee, Yee Thung Lee, Courtney Goodridge, Nadia Voigt da Mata, Ibrahim Ozturk, Peter Woodthorpe, Rafael Goncalves, Yue Yang, Yang Li, Stefanie Horn, Vishnu Radhakrishnan, and many others – thank you. Your passion for research inspired me, and I will never forget the family-like and teamwork-oriented environment of the HF&S research group.

To my friends and fellow PGRs at the Institute for Transport Studies (ITS) – Maximiliano Lizana Maldonado, Alexandra Elena Vitel, Vanessa Ternes, Victor Cantillo Garcia, Sazali Abdul, Ejiro Ikoko, Nur Karim, and many more – I will always cherish the nice times, the smiles and sharing our experiences.

To my friends back in Greece and across the world, thank you for your encouragement and for checking in on me.

Finally, to the people who are my entire world. To my family – my mother, Sofia, my father, Panagiotis, my brother, Gabriel, and my grandparents – thank you for making me the person I am today. Your unconditional love and support have been immense throughout my journey. And to my wonderful girlfriend, Katherine, thank you will never be enough. Your unwavering faith in me, your perspective when it was most needed, and your love at all times have been my greatest sources of courage.

Abstract

Understanding and modelling the interactions between pedestrians and automated vehicles (AVs) is important for facilitating widespread and safe AV deployment. During pedestrian-vehicle interactions, pedestrians form and update their beliefs regarding the vehicle's behaviour as it approaches. The mechanisms determining how a pedestrian interprets the behaviour of an approaching vehicle remain unclear. Previous studies have proposed models of cognitive mechanisms, such as estimating the goals of other agents, but none has attempted to apply and model the behaviour estimation mechanism in the real and dynamic context of pedestrian-vehicle interaction. Drawing inspiration from cognitive science, a first vehicle behaviour estimation experiment was conducted, and existing Bayesian observer models of goal estimation were modified and applied to the pedestrian crossing setting. Thus, an observation-based model with two alternatives based on 1) direct deceleration perception and 2) a more plausible visual cue, the rate of change of the relative rate of optical expansion $\dot{\tau}$, was proposed. The first experiment demonstrated that pedestrians do not solely rely on deceleration-related cues to judge whether an approaching vehicle is stopping, but that the vehicle's kinematic conditions, specifically its speed, time-to-arrival, and overall manoeuvre time history, also influence their beliefs. Even though the observation-based model achieved a relatively high correlation between model predictions and average pedestrian beliefs, it did not predict all the average pedestrian belief patterns in detail, being quite limited in predicting beliefs when the vehicle maintains constant speed. So, it was assumed that pedestrians may be utilising prior knowledge and situational expectations when the vehicle is far away, while deceleration observations become more crucial as the vehicle approaches. Thus, pedestrians likely infer the vehicle's behaviour by both directly observing the vehicle's actions and expecting the driver/AV to follow the most beneficial (value-maximising) behaviour. This rational, value-maximising reasoning mechanism was proposed as the value-based model. Both observation-based (Ob) and value-based (Vb) models were then integrated into an augmented model (Ob+Vb). All three models were evaluated for their ability to predict average pedestrian beliefs regarding the approaching vehicle's behaviour. This evaluation illustrated that Ob struggled with constant speed scenarios due

to its reliance on deceleration cues, Vb captured most kinematic effects but had limitations when the approaching vehicle is close to stopping, and Ob+Vb leveraged the strengths of both previous models, accurately reflecting all kinematic effects and belief patterns, achieving near-perfect correlation and the lowest error. Finally, to validate the three behaviour estimation models and test their generalisability, a second vehicle behaviour estimation experiment was designed and conducted, and the models' predictive capabilities were evaluated on the resulting dataset. Analyses on this dataset demonstrated the replication of previous findings in identical kinematic scenarios, validating the models, particularly Ob+Vb, which accurately predicted pedestrian beliefs, again. Furthermore, Ob+Vb successfully generalised to unseen scenarios with varied speeds and new manoeuvres, showing its ability to predict beliefs in novel situations. Additionally, Ob+Vb again exhibited superior performance, obtaining near-perfect correlation and the lowest error compared to the other models. Together, these studies demonstrate that: 1) while Bayesian observation of behaviour may suffice for simple laboratory tasks, it falls short in real traffic contexts, 2) pedestrians assess approaching vehicle behaviour by combining observations of vehicle actions with expectations of the driver's most rational, value-maximising future actions and 3) the proposed augmented model successfully predicts pedestrian beliefs, reproducing findings quantitatively and qualitatively, illustrating generalisability, and providing a likely explanation of the mechanisms with which a pedestrian interprets the behaviour of an approaching vehicle. Overall, investigating and modelling behaviour estimation in the pedestrian-vehicle interaction setting and its underlying mechanisms, present a significant challenge. However, this thesis demonstrates that it is possible to gain deeper insights into how a pedestrian interprets an approaching vehicle's behaviour, by integrating different psychological theories. This thesis not only enhances the theoretical understanding but also offers practical implications for designing safer and more intuitive interactions between pedestrians and AVs.

Contents

List of Tables

List of Figures

1 Introduction

1.1 Background and motivation

It was not until the 1970s that walking gained research interest as a mode of transport (Hitchcock and Mitchell, 1984) even though it has been the most widespread and originally the primary means of human locomotion. Currently, walking accounts for the second most popular mode of transport (in terms of the number of trips per year), after driving cars, in England, as per the Department of Transport's National Travel Survey (Department for Transport, 2024). Its popularity stems not only from its convenience, minimal environmental impact, and health benefits, but also from its relative importance to access other modes of transport and foster social connections (Dytrt, 2023; Loukaitou-Sideris, 2020). However, in transport research, the term "walker" is rarely used, but rather the most common term is "pedestrian" (Wigan, 1995), borrowed from the Latin *pedester* which means "on foot". Despite the etymology of the word, people with mobility impairments, sensory deficits or cognitive impairments, who use mobility aids to navigate, are also considered pedestrians (Federal Highway Administration, 2013). According to the Cambridge Dictionary the word "pedestrian" means "a person who is walking, especially in an area where vehicles go" (Cambridge Dictionary, 2024). The second part of that definition implies how important the interaction between pedestrians and vehicles is.

While cars have remained the most popular mode of travel and the English households' ownership, of at least one car, increased 5%, since 2002 (Department for Transport, 2024), it is apparent that more interactions or even conflicts between pedestrians and cars will happen (Zhao et al., 2019). The facts that pedestrians do not have protective equipment available and are moving slower than vehicles, make them one of the most vulnerable road users (VRUs) (El Hamdani et al., 2020). This vulnerability translates to a significant number of fatalities and injuries sustained by pedestrians in road accidents annually (World Health Organization, 2023a). The overall annual global toll of road traffic accidents resulting in deaths is approximately 1.19 million people, out of which pedestrians account for the 21% of those fatalities (World Health Organization,

2023b). Even though the overall road traffic deaths have decreased by 5% over the last decade, the deaths amongst pedestrians and cyclists remained almost the same (World Health Organization, 2023b). In the UK, the past two decades the total number of pedestrian deaths and injuries has decreased from approximately 35000 to approximately 19000 (Department for Transport, 2023). Even though this reduction was substantial, the absolute number remained incredibly high. In the US, during the last decade, only the pedestrian deaths increased and reached their highest number (7522 people) in 41 years (Naumann, 2025). While vehicle speed, impairment due to alcohol consumption, suboptimal road infrastructure and inadequate visibility of pedestrians are the most common factors resulting in pedestrian deaths and injuries (World Health Organization, 2013), another cause of pedestrian crashes is misunderstandings and incorrect expectations during pedestrian-vehicle interactions (Habibovic and Davidsson, 2012). The criticality of this issue and consequently the need to increase road safety shows the importance of the research on such traffic interactions (Markkula et al., 2020). During interactions with approaching vehicles, pedestrians must correctly interpret the vehicle's behaviour, to make their decisions regarding whether to cross the road or not and reach their own goal safely. So, investigating and modelling how pedestrians infer the behaviour of an approaching vehicle, from observing its movements, is important for the field of pedestrian safety.

The focus on road traffic interaction research has been intensified due to the technological advancements in robotics, sensors and ever-growing computational capabilities (Department for Transport, 2015; Nissan Motor Corporation, 2013), the rise of advanced driver assistance systems and the potential for fully autonomous vehicles (SAE levels 4 and 5; SAE International, 2021). While these autonomous vehicles (AVs) have promised significant societal and economic benefits (Centre for Connected and Autonomous Vehicles, 2022; Forrest and Konca, 2007; Reimer, 2014), ensuring safe and smooth interaction with pedestrians remains a key challenge (Rasouli and Tsotsos, 2020; Brown and Laurier, 2017; Millard-Ball, 2018). There are some concerns regarding these vehicles. Overly cautious AV behaviour could lead to frustration and delays on the roads (Brown and Laurier, 2017; Millard-Ball, 2018; Nordhoff et al., 2025). Conversely, unexpected actions or ambiguous communication signals by AVs

could cause accidents due to pedestrians' confusion or misinterpretation (Alambeigi et al., 2020; Habibovic and Davidsson, 2012; Lee et al., 2024). These potential problems highlight the urgent need for a deeper understanding of how humans interact in traffic, encompassing both qualitative and quantitative research (Kalantari et al., 2023; Rasouli et al., 2017a; Camara et al., 2020b; Schneemann and Gohl, 2016; Kotseruba and Rasouli, 2023; Markkula et al., 2023, 2018; Markkula and Dogar, 2022; Sadigh et al., 2018; Schwarting et al., 2019; Yang et al., 2024). Interactions between pedestrians and manually driven vehicles have been researched widely during the past few decades and it has been clear that it is not an easy task to fully understand and model them (Rasouli et al., 2017b; Guéguen et al., 2015; Markkula and Dogar, 2022; Sun et al., 2002). In the case of AVs, the absence of a human driver makes the situation even more complicated as the pedestrian-AV interactions lack communication aspects which are available in the pedestrian-driver interactions, for example eye contact or hand gestures (Rasouli and Tsotsos, 2020; Dey and Terken, 2017; Markkula et al., 2020). However, communication in road traffic interactions involves being able to convey intentions through movement or "body-language" and explicit signalling (if needed), confirming that those were understood and understanding the intentions of other road users (Ackermann et al., 2018; Fuest et al., 2018; Schieben et al., 2019).

So, without a doubt, increased road safety is still needed, especially for pedestrians, who are the most vulnerable in traffic interactions. The interaction between a pedestrian who is about to cross the road, and an approaching vehicle is very common and safety critical. For years, AVs were seen as a promising solution to keep pedestrians safe from accidents. However, researchers have realised that figuring out how these cars should interact with people is a big challenge (Schieben et al., 2019). AVs need to know the safest way to behave around pedestrians, but pedestrians can be unpredictable, making it hard to understand how they will react and what these cars should do. Researching and more specifically modelling pedestrian crossing behaviour is important for making roads safer and developing safe and interaction-capable AVs. The work presented in this thesis revolves around modelling the mechanisms with which pedestrians form and update their beliefs regarding the behaviour of an approaching vehicle, which communicates its intent only by its movement (no

explicit signals) and hopes to shed light on the relation of those mechanisms with the overall road crossing decision making process in such interactions.

In the upcoming sections a review of past literature is provided, including the key causation factors which lead to accidents in pedestrian-vehicle interactions, observations of pedestrian road crossing behaviour when they interact with approaching vehicles, how they communicate with such vehicles, models of the road crossing task, and observations and models of how humans infer the intentions of an approaching vehicle.

1.2 Road accidents involving pedestrians

As stated in the previous section, VRUs and particularly pedestrians are those most likely to be harmed in traffic accidents (Department for Transport, 2024; National Highway Traffic Safety Administration, 2024). This primarily includes pedestrians and people on two wheels (like cyclists and motorcyclists) because they lack the outer protection of a vehicle (El Hamdani et al., 2020). Among this group, pedestrians and cyclists are less likely to cause harm to others, but on the other hand, heavier and faster vehicles, can pose a risk to them. Among vulnerable road users, some groups are at a higher risk of injury, including the elderly, people with disabilities, and children.

As stated before, every year, approximately 1.19 million people are losing their lives in road traffic accidents around the world and VRUs accounted for more than half of those deaths (World Health Organization, 2023a). More specifically, in the EU, almost 20% of all road deaths involved pedestrians, which was a higher proportion in comparison to other VRU groups (European Road Safety Observatory, 2023). Similarly, in the US the percentage of pedestrian deaths is approximately 5% of all the road traffic deaths (National Highway Traffic Safety Administration, 2024).

Several risk factors play an important role in the likelihood and severity of pedestrian vehicle collisions. A major risk factor that increases the likelihood of a fatal injury for a pedestrian when struck by a motor vehicle is impact speed (Pikūnas et al., 2004; Tefft, 2013). Alcohol consumption is also a major risk factor for pedestrian accidents as it impairs judgment, slows reactions, and reduces

vision, making it more likely for pedestrians to be involved in accidents, in which either the driver (Phillips and Brewer, 2011) or the pedestrian (Lasota et al., 2020) is intoxicated. Absent or poorly designed or maintained road infrastructure (such as sidewalks, crosswalks, traffic signals) can lead to serious accidents as well (Zegeer and Bushell, 2006). Pedestrians are more likely to be hit if they are hard to see, which can be caused by lack of streetlights or if they are not wearing reflective clothing in low-light conditions (World Health Organization, 2004). Other factors that could increase the risk of pedestrian accidents include inattentive and distracted pedestrians (Ropaka et al., 2020). For the abovementioned issues, the solutions that have been most often suggested regarding increasing pedestrians' safety and preventing fatal collisions include interventions on the road design and infrastructure, the use of protective and alerting equipment and education around road safety (Yannis et al., 2020; Zegeer and Bushell, 2006).

Although the causation factors which are listed above are more common, they could be mitigated through safety interventions and prevention measures, as stated above (World Health Organization, 2013). Conversely, there are other causation factors for which it is not as clear whether prevention measures would help. A study by Thomas et al. (2013), which was based on an in-depth analysis of investigated crashes, identified key human-related factors. For both pedestrians and drivers, temporary person related factors (inattention/distraction due to competing external or internal activities), were the most common causation factor (25% and 28% respectively). A close second causation factor was found to be interpretation errors (16% for car drivers and 21% for pedestrians), which occurred either when the road users' expectations of what the other road users would do did not match with the reality or due to time/distance misjudgements. Most analyses/models regarding causation of crashes and near critical situations, in VRU/pedestrian-vehicle interactions, focus on the driver's perspective. The most common causation patterns from the driver's perspective have been found to be: unawareness of the conflict pedestrian (obstructed line of sight or reduced visibility), distracted driving, unexpected pedestrian behaviour, improper estimation of the gap distance between the vehicle and the pedestrian and infrastructure issues for example lack of pedestrian crossing (Habibovic et al., 2013; Sheykhfard and Haghighi, 2018; Yue et al., 2020). Habibovic and Davidsson (2012) have provided an exception to the driver-centric view, in their

analysis of accident causation in VRU-vehicle interactions, where they found that pedestrians' misinterpretations of the driver's intentions is a significant causation factor. They applied the SafetyNet Accident Causation System (SNACS) to a 995 crashes dataset, to systematically classify the crash causation information and organise that information into causation patterns. SNACS is a standardised method for systematically classifying crash causation factors. The findings of that study highlighted that in 70% of 56 crashes which involved VRUs (20 pedestrians and 36 bicyclists), the VRU was aware of the approaching vehicle but misinterpreted the situation and/or made an avoidance planning error. For example, they incorrectly thought the driver had noticed them and would react, or they miscalculated when to cross. Similarly, Räsänen and Summala (1998) suggested two main mechanisms producing bicycle-vehicle collisions, one being the driver's misdirected attention and the other one being the cyclist's erroneous beliefs about the driver's intentions. In the second case, those faulty assumptions were based either on law requirements or the preconception that the driver should be the one adjusting their behaviour to the behaviour of the more vulnerable road user, i.e. the cyclist. Thus, to make a step towards increasing pedestrian safety, there is a need of further research and especially modelling the cognitive mechanisms through which pedestrians infer the behaviour of an approaching conflict vehicle. To build such models, the first step would be to explore the overall pedestrian behaviour and the factors that influence their behaviour during interactions with approaching vehicles.

1.3 Pedestrian behaviour during road crossing interactions

Crossing the road while oncoming traffic is present is a basic and frequent task for all pedestrians. Unsurprisingly, most pedestrian crashes occur when people attempt to cross the road while there are vehicles approaching, or at crosswalks that are not very safe (Malenje et al., 2018; Sucha et al., 2017; Zhuang and Wu, 2011). Thus, studying pedestrian road crossing behaviour is important for improving road safety. Researchers have conducted extensive studies for years to gain insights into pedestrian crossing patterns and identify strategies for improving road safety. In addition to improving road safety researching pedestrian behaviour is also relevant to the development,

acceptance and eventually deployment of socially capable and safe AVs. Crossing the road is easy when there are no other road users, but it becomes a complicated task when there is oncoming traffic (usually vehicles). This more complicated situation can be regarded as an interaction between a pedestrian and an approaching vehicle and was chosen as the main traffic scenario in which this thesis was focused on. Markkula et al. (2020) defined road traffic interactions as "*situations where the behaviour of at least two road users can be interpreted as being influenced by the possibility that they are both intending to occupy the same region of space at the same time in the near future*".

Before modelling pedestrian road crossing behaviour during interactions with vehicles, it is important to review the key findings of past empirical studies on pedestrian behaviour. Firstly, the research area of pedestrian behaviour can be divided into macroscopic and microscopic (Papadimitriou et al., 2009). In the former case the pedestrian behaviour is approached in a more strategic level, for example overall planning of a journey, scheduling activities/goals and choosing departure times, whereas in the latter case the pedestrian behaviour is approached in an operational, local level, such as the road crossing task, interactions with other road users and obstacle avoidance. It has been argued that the macroscopic pedestrian behaviour research is lacking detail on the traffic interaction and is not that directly related to the safety aspect of such interactions (Papadimitriou et al., 2009). For these reasons the research done was only focused on the microscopic approach of pedestrian behaviour in this project and more specifically on the pedestrian road crossing behaviour.

Some typical examples of pedestrian road crossing behaviour are presented. Pedestrians frequently employ signals or actions to convey their crossing intentions to approaching vehicles, such as advancing forward, stepping onto the roadway, leaning forward, placing a foot on the road, observing oncoming traffic, or utilising informal signals (Risto et al., 2017). Rasouli et al. (2017b) found that the two most frequent patterns of pedestrian crossing behaviour, "standing, looking, crossing" and "crossing, looking", only account for half of their observed cases. Additionally, they highlighted that one-third of pedestrians in the non-crossing scenarios found to be waiting at the curb and looking at the traffic. These findings indicate high variability in the behaviours of pedestrians at the point of crossing/no-crossing. So, even though pedestrians

usually tend to stand on the sidewalk to signal their intention to cross, it is not apparently necessary that crossing the road is their goal (Crowley-Koch et al., 2011). In some cases, despite pedestrians initiating their crossing, showing their intention or the presence of markings and traffic signs, drivers do not always slow down enough as they approach crosswalks (Risto et al., 2017; Várhelyi, 1998). Sucha et al. (2017) highlighted that drivers were not very willing to yield to pedestrians who were just waiting at the curb and looking at oncoming traffic, with no obvious intention to cross. However, the distance between the pedestrian and the curb is a significant factor in influencing driver behaviour and can be used to predict the likelihood of a driver yielding (Himanen and Kulmala, 1988; Schroeder, 2008). Research has suggested that drivers are less likely to yield if a pedestrian is waiting more than half a meter from the curb (Sucha et al., 2017). Pedestrians cross the street in different ways based on factors like vehicle speed, available gaps and the number of lanes (Chandra et al., 2014; Lobjois and Cavallo, 2007; Pawar and Patil, 2015; Schmidt and Färber, 2009; Yannis et al., 2013). If the gap is not wide enough, pedestrians might step backwards or run to avoid being hit by a vehicle (Zhuang and Wu, 2011). However, in the cases where the approaching vehicle is yielding, most pedestrians wait until the vehicle comes to a complete stop before crossing, rather than relying on their own judgment of whether it is safe (Sucha et al., 2017). This has been further characterised by a bimodal distribution in pedestrians' road crossing decisions (Giles et al., 2019; Lee et al., 2022; Pekkanen et al., 2022; Tian et al., 2025, 2023). According to this observation, a proportion of pedestrians might choose to cross shortly after they perceive the approaching vehicle given that it would be safe to cross, and the rest of the pedestrians would wait for the vehicle to stop or would initiate crossing before it stopped (and it was obvious that it was coming to a complete stop before them).

It is important to understand in more detail when pedestrians feel safe to cross the street and what factors influence their decisions. There has been a large number of studies in the past that investigated the factors that affect the pedestrians' road crossing behaviour (Rasouli and Tsotsos, 2020; Bazilinskyy et al., 2019; Ezzati Amini et al., 2019; Ishaque and Noland, 2008). The factors influencing pedestrian behaviour can be divided into internal (sociodemographic and psychology factors) and external (environmental and traffic-specific factors).

1.3.1 External factors influencing pedestrian road crossing behaviour

Pedestrians' crossing decisions are influenced by a variety of external factors, which are primarily related to the behaviour (e.g., speed) and characteristics of the approaching vehicle (e.g. size), and environmental characteristics (e.g., weather/lighting conditions, infrastructure, and ways of communication – the latter is described in detail in the next section), and their interactions.

Road crossing decisions are based on how safe pedestrians feel (Brill et al., 2024) to accept an available gap between them and an approaching vehicle. To that end, they must judge whether the time and distance between them and the vehicle is sufficient to cross safely (Beggiato et al., 2017; Brewer et al., 2006; Dipietro et al., 1970; Harrell and Bereska, 1992; Moore, 1953; Nuñez Velasco et al., 2019; Petzoldt, 2014; Palmeiro et al., 2018; Schmidt and Färber, 2009; Wang et al., 2010). Time gap is defined as the temporal separation between two consecutive vehicles approaching the position of the pedestrian. The pedestrian can either accept or reject that gap. The minimum gap that a pedestrian would accept, in order to cross the road is called "critical gap" and is influenced by the vehicles' speed and distance, the crossing length, the pedestrians' speed and crossing initiation time, the pedestrians' characteristics and the road conditions (Department of Transport and Main Roads, 2006; Pawar and Patil, 2016). A recurring theme in the literature has been the relationship between the distance and time in pedestrian's road crossing decisions. While the time-to-arrival (TTA) has been considered as the most direct measure of safety (Petzoldt, 2014; Pugliese et al., 2020), studies have suggested that the vehicle's distance is what pedestrians might actually rely on in reality (Oxley et al., 2005; Schmidt and Färber, 2009; Yannis et al., 2020). This could lead to unsafe situations in which for faster moving vehicles, pedestrians accept shorter gaps, since for a given TTA the faster the vehicle, the further away it is from the pedestrian (Lobjois and Cavallo, 2007; Tian et al., 2022). Researchers have suggested that this could be explained by the human perception mechanisms. Specifically, Petzoldt (2014) showed that speed consistently affects the TTA estimates of pedestrians, i.e., their estimations of TTA were lower for low speeds and higher for high speeds. Hence, at higher speeds, they made risky crossing decisions because of their overestimations of TTA rather than because they used an incorrect decision-

making strategy. It was later found that faster vehicles produce weaker looming signals on the retinas of the pedestrians for the same TTA (Tian et al., 2022). Further support to that was the findings of Wang et al. (2025), who replicated such speed-dependent gap acceptance by assuming that road crossing decisions are boundedly optimal (noisy visual perception). Research has provided evidence about pedestrian's consistently rejecting very short gaps. Studies found that all pedestrians rejected gaps of less than 3 s (Dipietro et al., 1970; Schmidt and Färber, 2009). One of the earliest works in gap acceptance suggested that the minimum accepted gap would not be shorter than 2.5 s (Cohen et al., 1955). More recent studies highlighted that only a very small proportion of pedestrians accepted gaps shorter than 2 s (Lee et al., 2022; Pekkanen et al., 2022). The range of 3 to 6 s has been found to be the critical window, since the decisions within it are varied. Even though the gap acceptance for the 3 s is not as common, the rate is notable (Lee et al., 2022). As for the 4 and 5 s, research has found that the proportion of accepted gaps is around 50% (Cohen et al., 1955; Lee et al., 2022; Pawar and Patil, 2015). For gaps that are longer than 6 s, almost all pedestrians accept them. The probability of crossing the road for such gaps has been found to be very close to, if not 100% (Schmidt and Färber, 2009; Yannis et al., 2013). Eventually, pedestrians who are slower and avoid to take risks usually take longer to accept a gap (possible explanations are provided in the next subsection), leading to longer waiting times to accept a gap (Sun et al., 2002). After waiting too long for the critical gap, such pedestrians may accept shorter and riskier gaps (Antić et al., 2016).

The speed of a vehicle is one of the main factors that pedestrians consider when deciding whether to cross or not (Ackermann et al., 2018; Jiang et al., 2011; Pawar and Patil, 2015; Petzoldt, 2014; Sucha et al., 2017). As described before, the most frequent counter intuitive finding is the speed-induced unsafe crossing behaviour, where pedestrians tend to accept shorter time gaps when faced with faster vehicles. More specifically, in Petzoldt's (2014) study, the mean accepted time gap was found to be smaller for vehicles at 50 km/h (2.98 s) than at 30 km/h (3.57 s). Similarly, Tian et al. (2022) found that the gap acceptance percentage was higher for higher vehicle speeds and for a given time gap. This counter intuitive observation is illustrated in Figure 1.1 and is explained by the perceptual mechanisms discussed above.

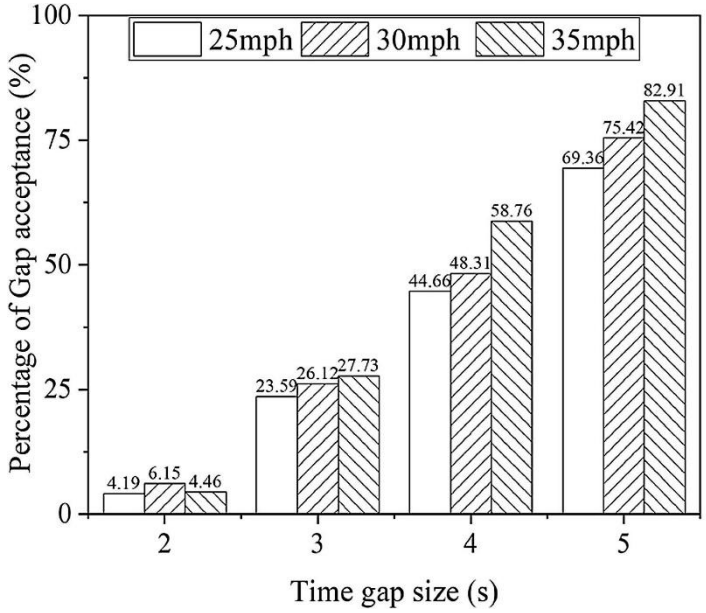

Figure 1.1 - Percentage of gap acceptance versus time gap size for different vehicle speeds. Figure adapted from Tian et al. (2022)

Vehicle speed also affects the behaviour of vehicles towards pedestrians and more importantly whether and how they would yield, with lower speeds leading to more yielding (Ackermann et al., 2018; Geruschat and Hassan, 2005; Himanen and Kulmala, 1988; Sucha et al., 2017; Turner et al., 2006). If drivers do not slow down as they approach a crosswalk, it can signal that they intend to maintain priority and not yield (Sucha et al., 2017). However, slowing down before reaching a crosswalk is important for pedestrian safety (Risser, 1985). Drivers need enough distance to react to pedestrians who suddenly appear (Várhelyi, 1998). Drivers may also need to slow down or stop completely to avoid hitting pedestrians who are distracted, running, or crossing unexpectedly/illegally (Katz et al., 1975; Sucha et al., 2017).

While spatiotemporal distance and vehicle speed are the primary external factors that affect the pedestrian road crossing behaviour, the (negative) rate of change of the vehicle's speed – its deceleration – is also an important factor, which influences the pedestrians' crossing decision. The deceleration has two key components: the timing (onset of braking) and the magnitude (deceleration rate). An early and gentle braking initiated at a distance from the pedestrian could be perceived as a clear indication that the vehicle is slowing and giving way, encouraging an earlier crossing (Dietrich et al., 2020; Risto et al., 2017; Tian et al., 2023). Conversely, late and harsh braking has the opposite effect and is perceived as ambiguous. Even though pedestrians can detect higher

deceleration rates faster, Ackermann et al. (2019) found that they detected decelerating behaviour faster at lower vehicle speeds. This could suggest that pedestrians might have the tendency to interpret lower speeds as indicative of the vehicle yielding even if it is not (Tian et al., 2023). Based on the different time gap durations, which were described above, it would be safe to assume that pedestrians expect to perceive vehicle yielding (deceleration) information between time gaps of 2 and 6 s. This assumption aligns with the observation from Schneemann and Gohl (2016) that drivers would start to initiate braking typically at 3 to 4 s from the position of the pedestrians, and the latest at 2.5 s. Beyond the vehicle kinematics related factors, other external factors affect pedestrians road crossing behaviour too.

Road features like the width of the road, the number of lanes, and pedestrian crossing facilities can affect how pedestrians and drivers interact, in terms of expectations of other users' behaviour and perceived safety (Brewer et al., 2006; Chandra et al., 2014; Ishaque and Noland, 2008; Lin et al., 2019; Pawar and Patil, 2015; Schroeder, 2008; Sucha et al., 2017; Turner et al., 2006; Zhao et al., 2019; Zhuang and Wu, 2011). A study found that the number of lanes a pedestrian needs to cross, can help predict whether drivers will yield (Turner et al., 2006). The type of pedestrian crossing facility can also affect the proportion of vehicles yielding. Studies have shown that drivers are more likely to yield at marked unsignalised crossings than at crossings with traffic management systems or engineering treatments (Schroeder, 2008; Turner et al., 2006). Pedestrians waiting on central refuge islands may accept shorter gaps than those waiting at the curb (Hamed, 2001), and they may also accept shorter gaps in narrow streets (Schmidt and Färber, 2009). The latter could be due to the fact that the oncoming vehicles are more reluctant to drive in such roads with high speeds and hence the yielding behaviour is supported (Fitzpatrick et al., 2007; Sucha et al., 2017; Zegeer and Bushell, 2006). Pedestrian refuge islands and clear markings can also improve driver compliance on roads with lower speed limits (Turner et al., 2006). However, Himanen and Kulmala (1988) found that road width and the presence of refuge islands did not significantly affect the driver's nor pedestrian's behaviour.

A lot more external factors have been studied regarding their effects on the pedestrian crossing behaviour, such as the size of the vehicle (Ackermann et al.,

2018; Hamed, 2001; Himanen and Kulmala, 1988; Tian et al., 2020; Pawar and Patil, 2015; Petzoldt, 2016; Sun et al., 2002), weather and lightning conditions (Ackermann et al., 2018; Harrell, 1991, 1993; Whetsel Borzendowski et al., 2013), vehicle deceleration patterns (Ackermann et al., 2018; Carlowitz et al., 2024; Dey et al., 2021; Lee et al., 2019b, 2024; Tian et al., 2023), pedestrian's own walking speed (Federal Highway Administration, 2013; Lobjois and Cavallo, 2007; Willis et al., 2004), waiting and observation time (Tian et al., 2024; Lobjois et al., 2013; Schmidt and Färber, 2009), traffic density and different means of communication between the pedestrian and the approaching vehicle (Rasouli et al., 2017b; Rasouli and Tsotsos, 2020; Chandra et al., 2014; Ezzati Amini et al., 2019; Tian et al., 2024; Lobjois et al., 2013; Pawar and Patil, 2015; Risto et al., 2017; Sucha et al., 2017).

1.3.2 Internal factors influencing pedestrian road crossing behaviour

Demographics, such as age and gender have been heavily investigated in relation to pedestrian road crossing behaviour. Generally, age has been found to be the most important factor which influences the cautiousness of pedestrians waiting to cross, with older pedestrians being more risk averse and cautious than younger pedestrians (Beggiato et al., 2017; Chandrapp et al., 2016; Harrell, 1991). This has been supported by Lobjois and Cavallo (2007), who examined the gap acceptance of older and younger individuals under different time constraints, finding that when there were no time constraints, older individuals accepted larger gaps therefore demonstrating less risky decision-making. Older pedestrians have also been found to wait on the curb for a longer time before accepting a gap in a similar display of more cautious crossing behaviour (Wang et al., 2010). Moreover, when it comes to estimating the vehicle's time to arrival (TTA) – an important estimate used to make road crossing decisions (Beggiato et al., 2017; Petzoldt, 2014), young children (5-6 years) were found to estimate TTA based only on the distance from the vehicle, as compared to a) adults (18-54 years), who were able to integrate both distance and speed for TTA estimations, and b) older children (7-10 years), who were found to gradually be developing that ability (Hoffmann, 1994). According to Piaget (1970), children's ability to correctly estimate time, distance and speed of objects is under development approximately up until the age of 8. However, later research has shown that children's poor judgments is not due to this inability, but due to the

problematic attention allocation on the presented task (Droit-Volet et al., 2006; Meir et al., 2013). So, the inadequate understanding of traffic rules, the lack of experience in the road crossing task, unpredictable behaviour, distraction and overly risky behaviour are common reasons behind child pedestrian accidents (Johansson et al., 2004; Leden et al., 2006; Rosenbloom et al., 2008). Conversely, Oxley et al. (2005), who studied pedestrian gap selection between 30-45, 60-69 and 75+ years old adults, found that all age groups used vehicle distance rather than TTA to make their decisions, however, the oldest pedestrians seemed to need more time to reach a decision. They also indicated that despite the generally riskier road crossing decisions, younger adults made safer choices as opposed to older pedestrians. Oxley et al. (2005) suggested that these unsafe decisions were a result of the deterioration of the older pedestrians' executive function (Staplin et al., 2001; Valos and Bennett, 2023), which includes features such as working memory, strategy application and cognitive flexibility (Gilbert and Burgess, 2008), rather than not being able to process the available perceptual cues. The greater risk aversion, possible misinterpretations of the situation and difficulties to make a decision could possibly be explained by the fact that older pedestrians (typically 60+) travel at lower speeds, 0.97 - 1.27 m/s, than younger adults, 1.32 – 1.57 m/s (Ishaque and Noland, 2008), and thus they generally select larger gaps to compensate for that lower walking speed. Observations by Bennett et al. (2002) regarding the lower crossing speeds of older (1.35 m/s) than other pedestrians (1.70 m/s) further support the previous findings. It is assumed that pedestrians that are most frequently present outside on the roads are young and middle-aged adults, as has been also indicated by the fact that these pedestrians have been comprising the majority or a large proportion of studies' samples (Beggiato et al., 2017; Harrell, 1991; Hoffmann, 1994; Oxley et al., 2005). Overall, middle-aged pedestrians have been found to not be as cautious as the elderly, and to accept smaller gaps (Beggiato et al., 2017; Harrell, 1991; Lobjois and Cavallo, 2007). Despite that, middle-aged pedestrians' crossing decisions could be described overall as safer than the decisions of children (Hoffmann, 1994) and the elderly (Lobjois and Cavallo, 2007), primarily due to their better performing executive function and movement ability, which result in better judgments of the traffic situation, and faster reaction and crossing times. Based on those better judgment skills, it has been suggested that middle-aged

adults are more aware of their surroundings when crossing the street and they tend to look at oncoming traffic more often (Zhuang and Wu, 2011).

Another important internal factor in the road crossing behaviour of pedestrians' is their gender. A lot of researchers have suggested that female pedestrians generally exhibit more cautious behaviour and risk aversion than male pedestrians. For example, studies have consistently shown that women have the tendency to wait for longer times at the curb before they decide to cross (Hamed, 2001; Tiwari et al., 2007). According to Holland and Hill (2007) and Himanen and Kulmala (1988), women have also been found to perceive risky crossing decisions to have a greater likelihood of resulting in a harmful outcome than men. Overall, these translate to higher cautiousness levels when deciding to cross the road for female pedestrians when compared to male pedestrians (Harrell, 1991). On the other hand, male pedestrians not only show a higher propensity in risk-taking (O'Dowd and Pollet, 2018), but also generally decide and cross the road faster than women (Lobjois and Cavallo, 2007). This difference between the road crossing behaviours of men and women might be explained by physiological and personality characteristics. Gender has been indicated to be the controlling factor of the relationship between impulsivity, physiological responses (such as skin conductance), and risk-taking, while at the same time conscientiousness is highly associated with safer pedestrian behaviour. Schiff and Oldak (1990) found that women underestimated the time to arrival more than, which could be another reason behind their more cautious road crossing behaviour. However, there have been studies which did not observe such gender differences on crossing decision tasks, such as time gap selection or safety margin calculation (Kadali and Vedagiri, 2016; Lobjois and Cavallo, 2007). Kadali and Vedagiri (2016) suggested that behavioural characteristics common to a big proportion of pedestrians might be nullifying the effect of gender in specific road crossing scenarios. The relationship between self-reported and observed road crossing behaviour is even more complex, for which Papadimitriou et al. (2016) found that on main roads the difference between reported and actual behaviour was not influenced by the pedestrians' gender but in minor residential roads women tended to overstate their risk-taking (reported crossing but were not), while men overstated theirs. The abovementioned findings suggest that the effect

of gender on road crossing behaviour is dependent also on the external factors as the ones mentioned in the previous subsection.

The size of pedestrian groups waiting to cross the street have been found to influence both pedestrian behaviour and driver behaviour. Larger groups are more likely to make drivers yield (Katz et al., 1975; Sucha et al., 2017; Sun et al., 2002). This can be safer for pedestrians as they would be more noticeable in groups (Zhuang and Wu, 2011). However, smaller groups can cause more traffic interruptions and delays due to frequent crossings (Jin et al., 2013; Malenje et al., 2018). Analysing illegal mid-block group crossings, researchers found that larger pedestrian groups require more time to cross. This can lead to drivers having to stop completely, potentially causing traffic waves. To avoid this, drivers may slow down in anticipation of pedestrians crossing, especially if there is a large group waiting (Malenje et al., 2018; Yi-Rong et al., 2015). Additionally, waiting times for pedestrians at zebra crossings can decrease as the group size increases (Hamed, 2001), possibly because a) drivers would often yield for a group of pedestrians about to cross the road in comparison to individual pedestrians (Katz et al., 1975), and b) pedestrians would utilise already established priority to pass by more pedestrians (Himanen and Kulmala, 1988).

There is plethora of other internal factors that are related to culture (Faria et al., 2010; Lee et al., 2011; Mihet, 2013; Sueur et al., 2013; Uono and Hietanen, 2015), familiarity of the place (Hamed, 2001; Sucha et al., 2017), social orientation (Evans and Norman, 1998; Harrell, 1993; Schwarting et al., 2019), psychological state (Berry and Schwebel, 2009; Cœugnet et al., 2019; Evans and Norman, 1998) and factors related to illegal behaviour (Jay et al., 2020; King et al., 2009; Pawar and Patil, 2015; Rosenbloom, 2009), which have been suggested to play a role in pedestrian road crossing behaviour.

After the presentation of the external and internal factors that influence pedestrians' road crossing behaviours, it can be acknowledged that there are several factors that affect the pedestrian road crossing behaviour. However, there are almost no empirical studies which focus on the question of how pedestrian understand the intentions of an approaching vehicle.

1.4 Communication between pedestrians and vehicles

As stated in the previous sections, communication between a pedestrian and an approaching vehicle is an important aspect of safe interactions and understanding it better is important for creating safe AVs. Markkula et al. (2020) have divided pedestrian-vehicle communication into two categories: communication through explicit and implicit cues. That means that pedestrians and vehicles communicate, by providing information to each other intentionally or even unintentionally. Explicit signals are actions that give information to others without changing the person's own movement or perception, for example generally conveying information verbally or a driver indicating that they are yielding to a pedestrian waiting on the curb by smiling or nodding to them (Rasouli et al., 2017b; Fuest et al., 2018; Mahadevan et al., 2018; Sucha et al., 2017). Some other common signals that are regarded as explicit cues are eye contact, hand gestures, and flashing headlights (Färber, 2016; Fuest et al., 2018; Sucha et al., 2017). Eye contact is important for pedestrians to make sure drivers are aware of them and sometimes to request the right of way. Hand gestures and light signals are less common and are used as last resort to resolve ambiguous situations or conflicts, or as an expression of gratitude or discomfort after the interaction (Färber, 2016).

Implicit signals are actions that change the person's own movement but at the same time these signals indicate the person’s intentions to others (Markkula et al., 2020). For example the body language of a pedestrian waiting at the sidewalk or a driver applying the vehicle’s brakes early or exaggerating their deceleration to show their intention to yield before a pedestrian (Fuest et al., 2018; Risto et al., 2017).

In the case of AVs, where a human driver is not necessarily present, pedestrian-driver communication is not possible. Past research has suggested that VRUs would find forms of explicit communication useful when interacting with AVs (Dey et al., 2021; Merat et al., 2018; Schieben et al., 2019). Examples include the acknowledgment that they have been detected by the AV or providing information about the AV’s status or intended behaviour (Carlowitz et al., 2023; Lee et al., 2019a). The main debate of most studies, regarding the behaviour of an AV, is whether external human-machine interfaces (eHMIs) are contributing to

the decision-making of the pedestrians or not. Some studies support eHMIs, as their findings suggest that they shorten the duration needed to recognise a yielding behaviour and increase the perceived safety and willingness to cross when it is needed (Böckle et al., 2017; Chang et al., 2017; de Clercq et al., 2019; Dey et al., 2021; Kooijman et al., 2019). Therefore, researchers and car manufacturers are exploring different designs for external interfaces that allow future autonomous vehicles to communicate effectively with pedestrians (Rasouli and Tsotsos, 2020; Bazilinskyy et al., 2019; Chang et al., 2017; Nissan Motor Corporation, 2015; Schieben et al., 2019). Researchers have developed different types of eHMIs for AVs to communicate with pedestrians. These interfaces often use light-bands, text messages, either on the vehicle or projected onto the road (Bazilinskyy et al., 2019; Lee et al., 2024; Tabone et al., 2021). Some also use auditory signals to convey the message (Deb et al., 2018; Lee et al., 2019a). However, different studies found different types of eHMIs more effective, and some have argued that already established ways of conveying a message like the flashing headlights are sufficient or even better than the current eHMI suggestions (Lee et al., 2022). Researchers have been studying how eHMIs affect pedestrian behaviour and feelings and the outcomes have been mixed. Some studies found that eHMIs can make pedestrians feel more comfortable and trusting of AVs, help make decisions more quickly when an AV had “eyes” on the pedestrian, reduce the time it took pedestrians to start crossing and can make pedestrians feel more positive about sharing the road with AVs (Chang et al., 2017; Deb et al., 2018; Dey et al., 2020; Holländer et al., 2019). Additionally, Dey et al. (2021) found that pedestrians were more likely to cross when an AV used an eHMI to signal that it was yielding. However, studies have shown that different eHMIs were not conveying any different message than the usual no-eHMI design to pedestrians (de Clercq et al., 2019; Deb et al., 2018; Kooijman et al., 2019), or provided messages that could be interpreted in different ways / be contradictory (Carlowitz et al., 2023; Lee et al., 2019a).

Additional evidence towards the importance of implicit cues is provided, as Domeyer et al. (2019) and Mahadevan et al. (2018) highlighted the importance of focusing on kinematics and not just the design of eHMIs. According to them, implicit communication or communication through actions is important and requires the ability to estimate the behaviour of others. Domeyer et al. (2022),

reported that in the case of non-intersection scenarios, the pedestrian acts accordingly to the vehicle's implicit cues. Rothenbücher et al. (2016), Clamann et al. (2017) and Moore et al. (2019), have agreed that pedestrians' decision making, when interacting with an AV, depends firstly on the vehicle kinematic cues and then on the eHMIs. For example, Schmidt and Färber (2009) suggested that distance is the cue affecting the pedestrian's decision, while they also studied the parameters that humans use to predict pedestrians' intentions. According to the abovementioned studies, implicit signals, alone, can impact how pedestrians make decisions when crossing the road, especially in situations where drivers do not have to yield to pedestrians at intersections. For instance, at uncontrolled intersections, pedestrians are more likely to cross the road if the distance and time to collision is shorter, and if the vehicle is traveling at a higher speed (Tian et al., 2022). The kinematic cues that have been found to have significant effect on the pedestrians' reaction times in detecting the deceleration of an approaching vehicle, and their willingness to cross (or not), and also on their crossing decisions, were the vehicle's deceleration rate and onset (Ackermann et al., 2019; Dey and Terken, 2017; Lee et al., 2024, 2022, 2019b; Petzoldt et al., 2018; Sucha et al., 2017; Wang et al., 2010; Várhelyi, 1998), and speed (Schneemann and Gohl, 2016; Petzoldt, 2016; Sucha et al., 2017; Várhelyi, 1998). Specifically, Dey et al. (2021) found that pedestrian willingness to cross the road increases when vehicles slow down dramatically. They also highlighted that pedestrian crossing willingness was not affected by eHMIs, but rather by the vehicle's kinematics in situations where the vehicle brakes aggressively. When approaching vehicles slow down early and brake gently, pedestrians are able to detect the vehicle's yielding intention more accurately and earlier and feel comfortable initiate crossing (Ackermann et al., 2019). However, if the braking is late and harsh, it leads to pedestrians' confusion and mistrust (Dey et al., 2021; Risto et al., 2017). To summarise, the above studies and their findings suggest that pedestrians find implicit signals as more trustworthy and consistent indicators of a vehicle's intention or behaviour than explicit signals, and that different deceleration patterns can have different effects on pedestrians' comprehension of the vehicle's exhibited behaviour.

Although a lot of research has been focused on the importance of overall communication in pedestrian-vehicle interactions and some of that on implicit

cues specifically, there is a lack of detailed empirical studies of how different vehicle manoeuvre patterns affect the mechanisms with which pedestrians estimate the intentions of approaching vehicles over the whole period of the pedestrian-vehicle interaction.

1.5 Modelling the road crossing task

It is valuable to complement empirical investigations with formal modelling, as it helps to build detailed theories, make more precise and focused inferences, explain and predict behaviours and improve the reproducibility of research findings (Guest and Martin, 2021), but also put the models to more direct use in applications. Computational models can help explain and reproduce the pedestrian road crossing behaviour and have real-world implications for traffic safety and infrastructure development (Markkula et al., 2023). The development of AVs has increased interest in this research area (Rasouli and Tsotsos, 2020; Camara et al., 2020a; Camara et al., 2021). AVs that do not understand pedestrian behaviour and interact appropriately will not improve traffic efficiency and safety (Markkula et al., 2020; Millard-Ball, 2018). Having appropriate and good models of pedestrian crossing behaviour will assist with the development of capable AVs, meaning that computational models of pedestrian behaviour are important for the wider deployment of AVs. Many models have been developed to account for different aspects of the overall pedestrian behaviour, but this thesis was focused on the road crossing decision-making.

Pedestrian crossing behaviour requires the combination of different mental processes. Palmeiro et al. (2018), building on Endsley's (1995) model of situation awareness (SA), suggested that there may be three levels involved: perception, comprehension and projection, before a crossing decision is made and performed. As can be seen in Figure 1.2, the pedestrian receives the relevant environmental information and approaching vehicle’s signals through the perception system, then combines the perceptual cues with prior knowledge and expectations and projects the vehicle’s status and predicts the vehicle’s most likely actions in the near future, before eventually they finalise the crossing decision and follow the respective action. Based on the division mentioned above,

the following section will provide a summary of theories and models of pedestrian road crossing behaviour. In this project the focus remained on modelling the three levels of situation awareness, which correspond to what it is meant in this thesis by behaviour estimation, which is also known as action understanding/intent recognition in the literature.

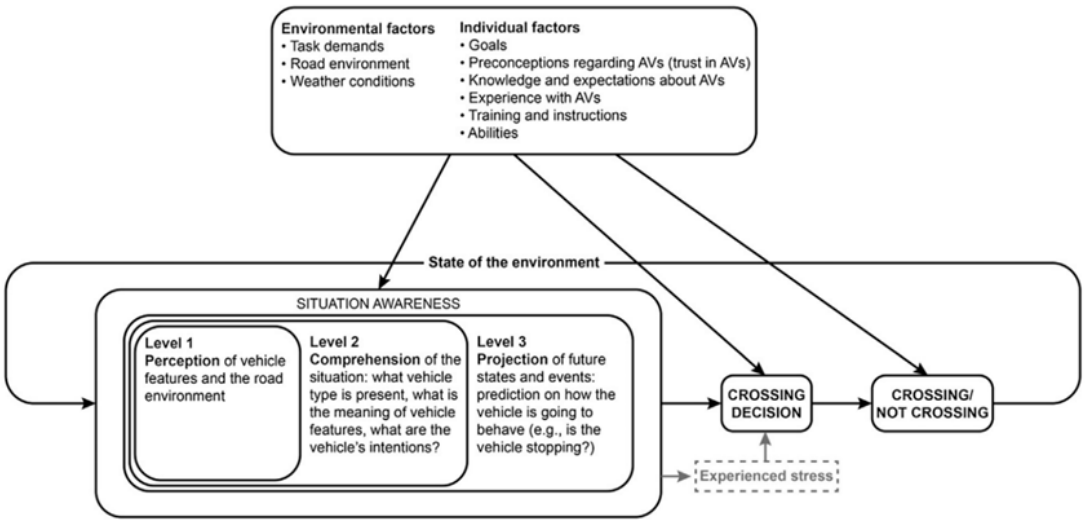

Figure 1.2 - Model of situation awareness (SA) in dynamic decision making, describing the interaction between a pedestrian and an AV – as illustrated by Palmeiro et al. (2018)[1] and adapted from Endsley (1995)

1.5.1 Perception theory for pedestrian road crossing modelling

It has been argued that, for pedestrians, comprehending the behaviour of approaching vehicle (yielding to the pedestrian or not) and updating their beliefs about such behaviour depend on mainly on implicit and sometimes on explicit cues as well (Ackermann et al., 2019; Lee et al., 2024; Petzoldt et al., 2018). As stated in Section 1.4, implicit signals play a significant role in pedestrian behaviour, but people do not seem to base their road crossing decisions on direct measurements of speed, time-to-arrival (TTA), distance, or deceleration rates (Lee et al., 2019b; Petzoldt et al., 2018). Instead, they rely on visual cues such as visual angle, the rate of change in visual angle, τ, i.e., the ratio of visual angle to the change rate of visual angle (could specify the instantaneous TTA of an approaching vehicle), and $\dot{\tau}$, i.e., the change rate of τ (DeLucia, 2015; Lee, 1976). When a vehicle approaches the pedestrians, the image of the vehicle on

[1] Reprinted from *Transportation Research Part F: Traffic Psychology and Behaviour*, Volume 58, Ana Rodríguez Palmeiro, Sander van der Kint, Luuk Vissers, Haneen Farah, Joost C.F. de Winter, Marjan Hagenzieker, "Interaction between pedestrians and automated vehicles: A Wizard of Oz experiment", p. 1006, Copyright (2018), with permission from Elsevier.

pedestrians' retinas increases as the distance of the vehicle to the pedestrians decreases (Lee, 1976). DeLucia (2015) suggested that pedestrians might perceive oncoming collision events from a particular angle, known as the bearing angle, which is the angle between the vehicle and the pedestrian's direction of movement. The expansion rate of the image is referred to as looming, which is a critical visual cue pedestrian may apply to judge events with a potential collision risk (Ackermann et al., 2019; DeLucia, 2015; Giles et al., 2019; Markkula et al., 2018; Tian et al., 2025). Generally, looming can be formulated as the change rate of the visual angle θ subtended by the vehicle, as follows: $\dot{\theta} = d\theta/dt$. When the vehicle approaches the pedestrians at a constant speed, looming increases as the distance decreases. However, in the vehicle yielding scenario, looming starts to decrease as the vehicle's speed decreases. Therefore, the change rate of looming may provide the vehicle's decelerating information for pedestrians. In other words, when the vehicle slows down, the edges of the vehicles seem to stop moving away from each other. Except looming, it is indicated that humans might apply a τ strategy to judge the collision events (DeLucia, 2015; Lee, 1976). Borrowing from the theory of visual control of braking, the optical variable τ is calculated by the ratio of θ and $\dot{\theta}$ ($\tau = \theta/\dot{\theta}$) and has been suggested to be an important visual cue, as it provides the time to collision information (Lee, 1976). It can be proven mathematically that the adequacy of the vehicle's deceleration to stop before a collision occurs is provided by the change rate of τ and $\dot{\tau}$. The collision will occur when $\dot{\tau}$ is less than -0.5, but not vice versa. The $\dot{\tau}$ is defined as the following equation: $\dot{\tau} = d\tau/dt$. In psychology, researchers have indicated that humans may apply several visual cues, e.g., distance, looming, τ, to detect collision events (Lee, 1976; Wann et al., 2011).

Pedestrian perception may depend on both visual cues and thinking strategies. Research has shown that pedestrians may estimate vehicle behaviour separately or as part of their decision-making process (Pekkanen et al., 2022; Tian et al., 2023). When there is a large gap in traffic, pedestrians may focus on the size of the gap rather than how the vehicle’s behaving (Tian et al., 2023). DeLucia (2015) found that when conflicts that can result to collisions are far away, people tend to use simple visual cues. However, the more imminent a collision event is, the more complex visual information are being used by the pedestrians. In the case of a road crossing interaction the perception level (Level 1 in Figure

1.2) could be characterised by the collision perception theory. So, collision perception theory is an important step to construct a pedestrian's situational awareness model.

1.5.2 Decision-making for pedestrian road crossing modelling

As discussed in Section 1.3, empirical observations show that pedestrians decide whether to cross the road by evaluating the size of the gap between them and the approaching vehicle. This has led to a group of models based on this assumption, called gap acceptance models (GA). Raff and Hart (1950) estimated the critical gap which is considered to be the threshold for a pedestrian to decide to cross. Later models built on this assumption (fixed critical gap models), where the critical gap is formulated as a function of factors like pedestrian speed or distance to the curb (HCM2010, 2010). Additionally, several recent models have estimated the critical gap by including vehicle speed and distance. Kotseruba's and Rasouli's (2023) model suggested that the critical gap decreases as the waiting time increases. A key assumption of the critical gap models is that all pedestrians are homogeneous. To avoid such limitation researchers introduced the binary logit models, where the crossing decisions are treated as a "yes/no" choice (Himanen and Kulmala, 1988; Sun et al., 2002; Zhao et al., 2019). Tian et al. (2022) combined a gap acceptance with psychophysics (based on visual cues) being able to capture patterns from pedestrian gap acceptance during continuous traffic flow. These binary logit models have been used with machine learning algorithms such as artificial neural networks (ANN), linear regression (LR) and support vector machines (SVM) to predict the pedestrians' crossing decisions based on individual and/or situational characteristics with great efficiency (Himanen and Kulmala, 1988; Pawar et al., 2016; Raghuram Kadali et al., 2014). However, these models lack interpretability (Srinivasan et al., 2023; Markkula and Dogar, 2022; Rudin, 2019), because they learn complex patterns to make accurate predictions, but their internal operation is not necessarily mapping to underlying processes that drive human behaviour. Other researchers have split the road crossing decision-making into strategies depending on the phase of the pedestrian-vehicle interaction (Tian et al., 2025). All the above models are based on observed behaviour patterns and are rather descriptive, except for Tian's models, which combined descriptive methods with plausible visual information. On the other hand, there have been researchers who modelled the underlying

decision-making mechanisms based on psychology theories and that is beneficial because they are more explanatory and generalisable than the previous group of models. Thus, computational models that are based on psychological theories can expand to explain scenarios where the approaching vehicle is yielding or capture the time dependency of a road crossing.

The development of cognitive models has benefited the somewhat “rigid” models that were mentioned before. Similarly, Wang et al. (2023), Crosato et al. (2023) and Markkula et al. (2023) combined models from cognitive science theories with machine learning approaches in order to explain patterns of pedestrian crossing behaviour. Another group of models which are based on psychological and cognitive theories, are the evidence accumulation models (Giles et al., 2019; Markkula et al., 2018; Pekkanen et al., 2022). These models have shown that they can overcome the limitations that the gap-based models are struggling with. Building on the well-established drift-diffusion process theory in psychology and cognitive neuroscience, these evidence accumulation models propose that pedestrian crossing decisions are the result of a process where visual cues and noisy evidence are accumulated (Ratcliff et al., 2016). When the accumulated evidence reaches a certain threshold, a decision is made. The resulting response time distribution provides insights into crossing decisions and their timing. These models offer a powerful tool for explaining pedestrian crossing decisions guided by perceptual cues from a human cognitive perspective. Other notable models that have been used to model the dynamic road crossing decisions are based on game theoretical approaches (Kalantari et al., 2023; Camara et al., 2021; Wu et al., 2019). An exensive table of details regarding the different models, their inputs and outputs presented in this subsection, can be found in Table E.1 of Appendix E .

1.5.3 Action for pedestrian road crossing modelling

As stated in the previous subsection, decision-making models of pedestrians can capture the time dependency of the road crossing task. Through them, the crossing initiation time (CIT) – the time it takes for a pedestrian to start crossing the road, can be calculated. CIT is affected by both internal and external factors (see section 1.3). For example, males’ and younger adults’ CIT is higher than that of females and older adults respectively (Lobjois et al., 2013; Lobjois and Cavallo, 2009). Also, when a pedestrian is facing a faster vehicle their CIT increases (Tian

et al., 2022). Evidence accumulation models have been successful in capturing the bimodal distribution of road crossings in vehicle-yielding scenarios (Pekkanen et al., 2022). The distribution of late crossing initiation times is complex and cannot be described by standard response time distributions, while the distribution of the early crossing initiation times is similar to that in non-yielding scenarios, as pedestrians use similar decision-making strategies (Tian et al., 2023). That leads to the suggestion that when the approaching vehicle is yielding, the pedestrian might have to employ a road crossing strategy based on a behaviour estimation mechanism (more details in Section 1.6).

After pedestrians finalise their decision, they initiate crossing and walk across the road. Walking behaviour has been replicated in microsimulations of models like social forces model (SF), cellular automata models (CA) and learning-based models. SF models have been used to simulate pedestrian-vehicle interactions and large-scale pedestrian flows and are based on Newtonian physics (Helbing and Molnár, 1995; Moussaïd et al., 2010; Zeng et al., 2014). CA have been good models for simulating complex environments due to their discreet definition (Layegh et al., 2020; Lu et al., 2016). In contrast to the white-box models that were mentioned up until now, there are black-box models based on learning-based approaches. These approaches, learn pedestrian walking behaviour from data. These models use techniques like artificial neural networks (ANN) (Song et al., 2018; Ma et al., 2016), Long Short-Term Memory networks (LSTM) (Kalatian and Farooq, 2022), reinforcement learning (RL), and inverse reinforcement learning (IRL) (Crosato et al., 2023; Martinez-Gil et al., 2014; Nasernejad et al., 2023; Wang et al., 2023) to simulate and predict pedestrian movements. The input on these models can be either the outputs of other models, for example SF, or image/video datasets.

Despite the variety of models regarding the road crossing task, there has been little focus and modelling efforts on the mechanisms with which pedestrians infer the intentions of approaching vehicles (comprehension and projection levels of SA – see Figure 1.2), which seems to be quite important especially in more complicated scenarios, like the ones that include changes in the speed of the vehicle, which in turn are very common in real life.

1.6 Modelling behaviour estimation mechanisms

The previous sections highlighted that pedestrians rely primarily on the implicit kinematic cues of an approaching vehicle to communicate with it and to decide when to cross the road and provided an overview of models of the overall road crossing task. The current section presents information about how humans infer the intentions of others and more specifically how pedestrians' beliefs regarding the behaviour of an approaching vehicle are formed and updated over time. Thie cognitive mechanisms of inferring the vehicle's immediate behaviour (particularly whether it is stopping or not) is what was defined as "behaviour estimation" in this thesis.

1.6.1 Human intent prediction models for AVs

Several studies have concentrated on modelling the road crossing behaviour of pedestrians, as discussed in the previous section but also on developing systems that promise improved pedestrian safety on the streets. Predicting pedestrian actions on roads is a safety concern for AVs and has become increasingly important to the automotive industry. Estimating when pedestrians will cross streets has proven to be a difficult task, as they can move in various directions, change their movements unexpectedly, be obscured by obstacles, and become distracted while talking or using their phones (Ferguson et al., 2015). It is apparent that their decisions and overall behaviour can be affected by several factors, as detailed in previous sections. To develop AVs, a lot of effort has been spent on algorithms for pedestrian detection. The data that these detection algorithms use include images, 3D point clouds, or a combination of both (Ferguson et al., 2015; Gandhi and Trivedi, 2008; Schneider and Gavrila, 2013). Research on predicting pedestrian behaviour has been focused on both short-term and long-term applications. Long-term prediction studies often have utilised static cameras to predict either the final destination or the trajectory to be followed by pedestrians (Deo and Trivedi, 2017; Karasev et al., 2016; Kitani et al., 2012). However, long-term predictions have been very challenging to obtain due to the easiness with which pedestrians can decide to alter their movements (Ferguson et al., 2015; Gandhi and Trivedi, 2008). On the other hand, despite the challenges posed by rapid changes in pedestrians' movements, short-term approaches have been able predict pedestrian trajectories within horizons of a few seconds.

Rehder and Kloeden (2015) highlighted the relevance of head orientation and body movement for short-term predictions, while long-term predictions are more goal oriented. Researchers have used a pedestrian's posture, and body language to infer their intentions. The models that have been applied for this application are mostly data-driven, for example models based on neural networks (Hariyono and Jo, 2015), Gaussian Process (GP) models (Quintero et al., 2015; Quintero et al., 2014) and Support Vector Machine (SVM) models (Koehler, 2015; Koehler et al., 2012). Kalman Filters and Particle Filters have been used in many dynamics-based studies to predict the pedestrians' positions or paths (Hariyono et al., 2015; Bertozzi et al., 2004, 2004). Information about head orientation has been used to improve pedestrian intention estimation. Several studies have combined pedestrian dynamics and situational awareness to predict pedestrian intentions (Schulz and Stiefelhagen, 2015a, 2015b; Hashimoto et al., 2015; Hashimoto et al., 2015). Some notable approaches that have incorporated head orientation information to the overall dynamics are the Multilayer Perceptron (MLP) network (Goldhammer et al., 2015) and Latent Dynamic Conditional Random Fields (LDCRF) system (Schulz and Stiefelhagen, 2015a, 2015b). However, relying solely on head orientation may not be sufficient, as it may not always indicate the pedestrian's current real attentiveness. The latest advances on the intention estimation of pedestrian research have been made by including the influence of the environment and/or the relations amongst all the involved interacting road users. To achieve that researchers have used various techniques, including Recurrent Neural Networks (Bock et al., 2024), Gaussian Processes (Ferguson et al., 2015; Quintero et al., 2015; Quintero et al., 2014), and Dynamic Bayesian Networks (Hashimoto et al., 2015; Hashimoto et al., 2015; Kooij et al., 2014).

Research on VRUs' intentions estimation has drawn a lot of attention (Ahmed et al., 2019; Chen et al., 2023; Kwak et al., 2017; Ranga et al., 2020; Saleh et al., 2020, 2018a, 2018b, 2017a, 2017b, 2017c; Sharma et al., 2022) and it is a crucial component in the development of AVs. But since the behaviour of those humans will also depend on their understanding of how the AVs behave (Habibovic et al., 2018; Jayaraman et al., 2019; Razmi Rad et al., 2020), it is interesting to consider how humans themselves estimate intent of other humans and especially road users – this is an understudied area, which is also of relevance in conventional

traffic safety, since it is known that many crashes occur because of road users misinterpreting each other's intentions (Ljung Aust et al., 2012; Yue et al., 2020).

1.6.2 Behaviour estimation in cognitive science

Humans are quite good at inferring the goals, beliefs, and desires of others through observing their actions (Baker, 2012; Gergely et al., 1995; Woodward, 1998) and this has been the focus of much research in the field of psychology and cognitive neuroscience. It has been suggested that when someone is observing someone else's actions, the observer is able to infer the goals and intentions of the other, due to the so-called mirror neurons and the cortical regions of the action-observation network (Kilner, 2011). During interactions, humans tend to interpret the behaviours of others as goal-directed actions/intentions (Gergely et al., 1995; Gergely and Csibra, 2003). Human social interaction depends on the ability to estimate the psychological states that produce behaviour (Baker et al., 2009). Intent recognition/behaviour estimation has been well researched in terms of establishing models that could explain their mechanism in a general sense, though they are limited to tasks such as reach-to-grasp an object (Amoruso and Urgesi, 2016; Iacoboni et al., 2005) and listening to a birdsong (Friston and Frith, 2015; Friston and Kiebel, 2009a, 2009b; Friston and Frith, 2015). Additionally, numerous other studies have applied models of goal/intent estimation, but they all shared the limitation of observing an agent navigating in a simple maze (Baker et al., 2011, 2005, 2017, 2009; Ramírez and Geffner, 2010). In these studies, the participants were shown stimuli of animated agents (and their trajectories) moving towards goals in simple two-dimensional maze-like environments and were asked to report their beliefs about the agent's goal at pre-fixed judgment points (i.e., points in the middle of the agent's trajectory before a particular goal was achieved, where participants reported their subjective inferences regarding the agent's goal), where the agent's movement sequence was paused. The models proposed by the literature above have been incredibly successful at capturing the approximately rational inference mechanism in human goal inference. These models were based on the idea of the Theory of Mind (ToM) – humans' ability to reason about other people's mental states – and were formalised as a Bayesian (BToM) inversion of a probabilistic state-estimation and expected-utility-maximising planning process, conditioned on observing others' actions. This BToM framework followed the

principle of rationality – the expectation that others form an approximate optimal plan to achieve their goal and was formulated as a Markov decision process (MDP) model of goal-directed planning, where the posterior probability of a Goal (belief) was calculated, conditioned on observed Actions and the Environment, using Bayesian inference:

$$P(Goal|Actions, Environment)$$
$$\propto P(Actions|Goal, Environment)P(Goal|Environment)$$

Where, $P(Actions|Goal, Environment)$ is the likelihood of the Goal given observed Actions and the Environment, defined above as probabilistic planning in an MDP, $P(Goal|Environment)$ is the prior probability of the Goal given the Environment, which sets up a hypothesis space of goals that are realisable in the environment.

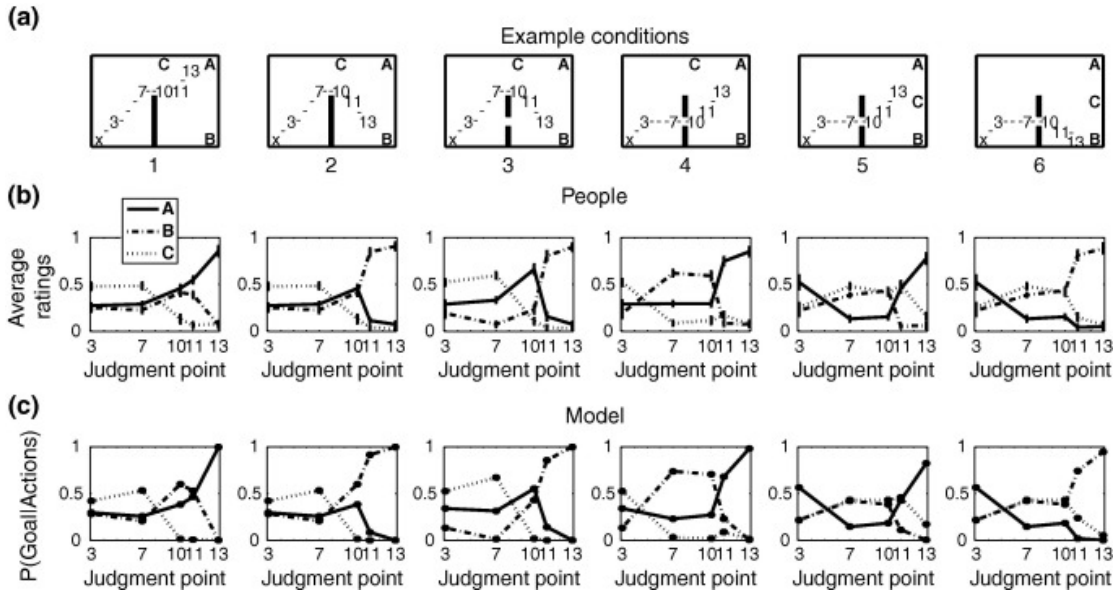

Figure 1.3 - Examples of qualitative comparisons between participants' goal inferences and the predictions of model M2 of Baker et al. (2009)[2]

An important illustration of their findings and methods can be seen in Figure 1.3. Panel (a) shows the experimental setup, where A, B and C represent the possible goals that the agent may pursue (with X being the starting point), and the numbered points indicate where the agent's movement trajectory was cut and a subjective judgment was given by the observer participant. Panels (b) and (c) compare the observed evolution of the average human beliefs with the predictions of the computational model, showing how the likelihood of each goal

[2] Reprinted from *Cognition*, Volume 113 /Issue 3, Chris L. Baker, Rebecca Saxe, Joshua B. Tenenbaum, "Action understanding as inverse planning", p. 337, Copyright (2009), with permission from Elsevier.

changes as new movement evidence is obtained over time. The average subjective ratings showed several patterns of reasoning: participants initially were uncertain regarding the agent's goal, but their beliefs became more certain as they accumulated more evidence from observing the agents' trajectories. Their beliefs were updated over time, and changed, for instance when an agent took a late turn towards a new goal.

Additionally, humans might be inferring others' goals not only based on observing the other agents' actions but also may be relying on their expectations about how a rational agent should behave (Jara-Ettinger et al., 2020; Markkula et al., 2023). This concept of rational, value-maximising reasoning aligns with affordance theory, which posits that perception is not just for interpreting the world, but for identifying potential actions and their value (Chemero, 2003; Lio et al., 2020). Therefore, a complete model of behaviour estimation may need to account for not only the observation of actions but also the expectations of the most rational or value-maximising actions.

1.6.3 Behaviour estimation in the traffic setting

Before a pedestrian decides to cross the road or wait for the upcoming vehicle to pass first, they might have to first perceive and estimate the intentions of the approaching vehicle, as was described in Figure 1.2. Pekkanen et al. (2022) and Tian et al. (2023) have provided quantitative proof and empirical findings, respectively, that support that pedestrian road crossing decisions involve a process of behaviour estimation, as described in the perception paragraph of Section 1.5. More specifically, Pekkanen et al. (2022) were able to account for the timing of late crossings by including the time derivative of TTA (i.e. $\dot{\tau}$, in the case that TTA is approximated by τ) as one of the accumulated decision sources of evidence, indicating that pedestrians may be using a process like deceleration estimation when deciding when to cross the road. In addition to the previous finding, Tian et al. (2023) found that when the time gap between the vehicle and the pedestrian is large, crossing decisions and pedestrians' judgments about the behaviour of the vehicle are negatively correlated, but when the vehicle is close and yielding, crossing decisions and pedestrians' judgments are positively correlated, indicating that pedestrians follow two different strategies to determine their crossing decisions: a) an early crossing decision based on a safe distance and TTA or b) a late crossing based on the speed and yielding behaviour of the

vehicle. Markkula et al. (2023) have combined a range of existing computational theories from different psychology subdisciplines into a joint model capable of capturing behavioural phenomena apparent in a variety of road traffic interactions. Part of those theories was also a theory of mind regarding others' intentions. Even though this theory of mind was part of a successful model as a whole, it was not tested in a behaviour estimation task specifically.

To sum up, much research has been done on machine learning (ML) models for estimating pedestrian intentions, but it seems that there is a lack of research on understanding and models of how pedestrians estimate the behaviour of an approaching vehicle. Also, there is a lack of investigation of the role of perceptually plausible visual cues in behaviour estimation. Finally, even though models of goal estimation have been successful in capturing humans' inferences when observing simulated agents in cognitive science laboratory paradigms, there is a lack of translation of those models to significantly more complex contexts, and more specifically the road traffic setting in this case.

1.7 Research gaps and objectives

Despite extensive research on pedestrian crossing decision modelling, several important questions remain unanswered. This thesis sought to address some of these questions by examining pedestrian road crossing behaviour and integrating these observations into computational models. The following critical gaps have been identified. Overall, whilst studies investigating the factors that affect the pedestrian road crossing task are extensive (Rasouli and Tsotsos, 2020; Ishaque and Noland, 2008), there are almost no empirical studies focusing specifically on how pedestrians estimate the behaviour of approaching vehicles. Concerning this gap, an in-depth and time-dynamic investigation of what are the factors affecting the underlying mechanisms of behaviour estimation from the pedestrians' perspective remains unexplored. Previous studies on pedestrian crossing behaviour have often focused on simple traffic scenarios with approaching vehicles moving at constant speed (Dey et al., 2021; Tian et al., 2024; Zhao et al., 2019). However, real-world traffic often involves vehicles with various acceleration patterns and changing behaviours, like a variety of yielding

manoeuvres (Ackermann et al., 2019; Lobjois et al., 2013; Risto et al., 2017). The underlying mechanisms with which pedestrians form and update their beliefs about the behaviour of an approaching vehicle in more complex traffic situations remain understudied. Misinterpreting a driver's intentions, such as incorrectly assuming a vehicle will yield or misjudging its approach, is a major causation factor in pedestrian accidents and can lead to fatal accidents for pedestrians. Additionally, for AVs to be safe and effective, their intended actions, as conveyed through their motion, must be unambiguously understood by pedestrians. While Bayesian Theory of Mind (e.g., Baker et al., 2009) and Value-Maximising Expectations (e.g., Markkula et al., 2023) offer possible theoretical explanations for how pedestrians form and update their beliefs, there has been insufficient experimental and modelling work exploring, in detail, the mechanisms of behaviour estimation from the pedestrian's perspective in the road crossing context. Building on the previous two research gaps and on the motivations to improve pedestrian safety and make interaction-capable AVs, a third gap naturally emerges. There is a lack of validated and generalisable models which provide explanations of the underlying mechanisms of behaviour estimation or in other words on how pedestrians form and update their beliefs regarding the behaviour of an approaching vehicle, that could cause a conflict, over time.

To address the research gaps mentioned above, this study had three objectives. It aimed to a) investigate the factors that affect the cognitive mechanisms by which pedestrians estimate vehicle behaviour, b) understand the underlying theories of such mechanisms, and c) create computational models of behaviour estimation in realistic road crossing situations. To achieve these objectives, this thesis implemented computational models derived from established psychological theories. Two experiments were designed to gather pedestrian belief data in a variety of traffic scenarios. The collected data were used to test three distinct models:

1) An Observation-based behaviour estimation component model (inspired by Bayesian observer models) that formulated behaviour estimation as a mechanism of updating beliefs based on deceleration-related kinematic observations.
2) A Value-based behaviour estimation component model that assumed pedestrians estimate behaviour by reasoning about the most rational,

value-maximising actions for the approaching vehicle's driver. While the simulated vehicle used in the experiments was a Level 5 AV with no visible driver, the model assumed that pedestrians interpret the AV's behaviour using the same mechanism that they apply to a human driver. This assumption was based on the fact that the current research was focused on the vehicle's implicit kinematic cues as the primary source of information. This could allow the Value-based component model to be applied to pedestrian interpretations of the behaviour of both AV and human-driven vehicles.

3) An augmented model that integrated the two component models of behaviour estimation mechanisms.

The structure of this research, outlining how these models were implemented and tested across the thesis, is depicted in Figure 1.4.

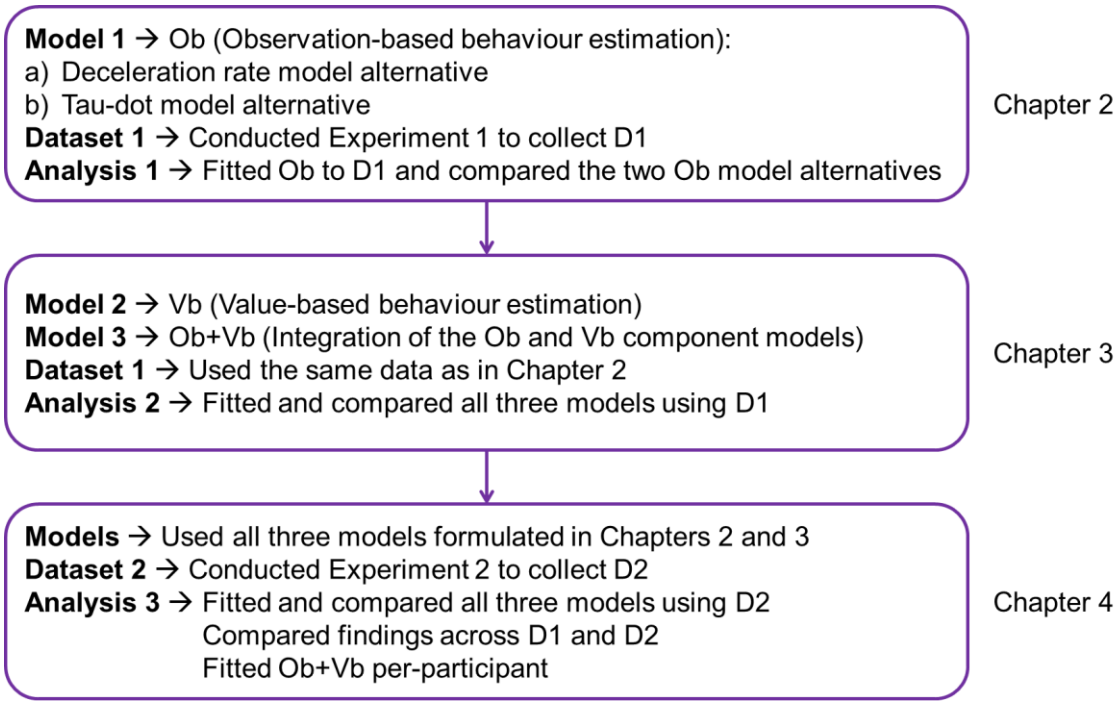

Figure 1.4 - Overview of the structure of the research

1.8 Thesis structure

This section outlines the structure of the thesis, briefly introducing each chapter to demonstrate how the research addressed the previously set objectives and contributed to filling the identified gaps in the literature.

Chapter 2, “Behaviour estimation through Bayesian observer models”, presents the foundation of the current research by detailing the design of the first experiment and introducing the first behaviour estimation component model, inspired by Bayesian observer models. This model is tested against the experimental data to evaluate the strengths and limitations of a purely observation-based mechanism of behaviour estimation.

It is important to note that the experiment described in this chapter was a collaborative effort with Dr Kai Tian, which formed the basis of a publication (Tian et al., 2023). As detailed in the “Intellectual Property and Publications” section, the candidate was responsible for the initial conceptualisation, experimental methodology, data collection and data curation alongside Dr Kai Tian. While the analysis in that publication and subsequently in Kai Tian's thesis (Tian, 2023) focused primarily on the empirical findings of a) the pedestrian crossing decisions, b) the pedestrian subjective judgments of the approaching vehicle’s behaviour and c) the combination of the two, this thesis uses the shared dataset for a different purpose, that is to implement, validate and generalise models of the underlying cognitive mechanisms of behaviour estimation.

Chapter 3, “Behaviour estimation through observations and expectations”, addresses the limitations found in the previous chapter by introducing a second behaviour estimation component model based on value-maximisation of rational agents. It then develops an augmented model that integrates both the observation and value-based mechanisms. All three models are fitted to the dataset collected in the experiment of the previous chapter and compared with a model selection technique.

Chapter 4, “Validating and generalising the behaviour estimation models”, tests the robustness of the behaviour estimation models, and more specifically the augmented one. A second, more comprehensive experiment with novel driving manoeuvres is used to validate the models against replicated scenarios

and investigate its ability to generalise to new situations. Again, all three models are fitted to the new dataset and compared with the same model selection technique as in the previous chapter. Finally, the augmented model is fitted on the per-participant level.

Finally, Chapter 5, “General discussion” provides an overview of the key and contributions of the research of the current thesis. It discusses the empirical, methodological, theoretical and practical implications of the research for road safety and automated vehicle design, acknowledges the study's limitations, and outlines promising directions for future work.

2 Behaviour estimation through Bayesian observer models

This chapter introduces an initial model of how pedestrians use visual observations to form and update their beliefs about an approaching vehicle's intent. The model focuses on scenarios where the vehicle communicates its intentions solely through its movement. As discussed in Section 1.6, numerous computational modelling attempts have used Bayesian observer approaches, in order to explain humans' inferences regarding another agent's goals, based on the latter's actions. These models have been evaluated by experiments using simplified stimuli. These stimuli were animated representations of agents (small moving circles and traces of their trajectories trailed behind them) navigating maze-like environments (a discreet grid of squares with walls displayed as solid black barriers), presented from an overhead perspective. While successful in these abstract contexts, the applicability of such Bayesian frameworks to the distinct and dynamic interactions between pedestrians and approaching vehicles remains unexplored. Therefore, this research adapts the Bayesian observer approach to model belief formation and updating in this specific pedestrian-vehicle setting. The proposed model is subsequently evaluated through an experiment. This experiment, while inspired by the paradigms used in earlier studies of goal estimation from observing the actions of another agent (Baker et al., 2009), is specifically designed to investigate how pedestrians' beliefs regarding vehicle intent evolve in scenarios that reflect key aspects of real-world road crossing situations.

2.1 Experiment

This section details the experimental methodology employed to investigate pedestrians' judgments regarding the behaviour of an approaching vehicle. An immersive virtual environment was used to present participants with various approaching vehicle scenarios involving different vehicle kinematics. Participants performed a road crossing task and a subsequent behaviour estimation task, where they judged whether an approaching vehicle was stopping or not. The data

collected from the behaviour estimation task were then used for statistical analysis (described in Section 2.2) and to develop and validate the computational model of belief updating (described in Section 2.3).

2.1.1 Participants

An experiment was conducted to investigate the research questions related to vehicle behaviour estimation. The study received ethical approval from the University of Leeds Ethics Committee (No. LTTRAN-145). 30 healthy adults, including 17 males and 13 females, aged between 20 and 67 (mean age = 30.73, standard deviation = 8.63) were recruited from the University of Leeds Virtuocity participant list. The participants were required to have no significant mobility issues or medical conditions such as epilepsy. They also needed to have either normal or corrected-to-normal vision and have lived in the UK within the last 12 months as their experience with road traffic could influence their road crossing behaviour and judgments. They provided written informed consent before participating and were given £15 as a reward for their participation.

2.1.2 Apparatus

The experiment was conducted at the Highly Immersive Kinematic Experimental Research (HIKER) lab at the University of Leeds. The pedestrian simulator is a CAVE-based simulated environment that utilises three glass wall projections and a floor projection, as illustrated in Figure 2.1. Participants were able to move in the simulated environment with a 9 m × 4 m walking space. The eight 4K projectors behind the glass walls or above the floor projected the scenarios at 120 Hz. Eight computers controlled the projectors and tracking system, which was data logging the participant's position through the tracking glasses they were wearing on their head, to adjust the projections in line with the participant's perspective. The virtual environment was created using the Unity3D software, where it is possible to record the kinematics information of the vehicles and participants, such as speed, position and experiment state, on each time step.

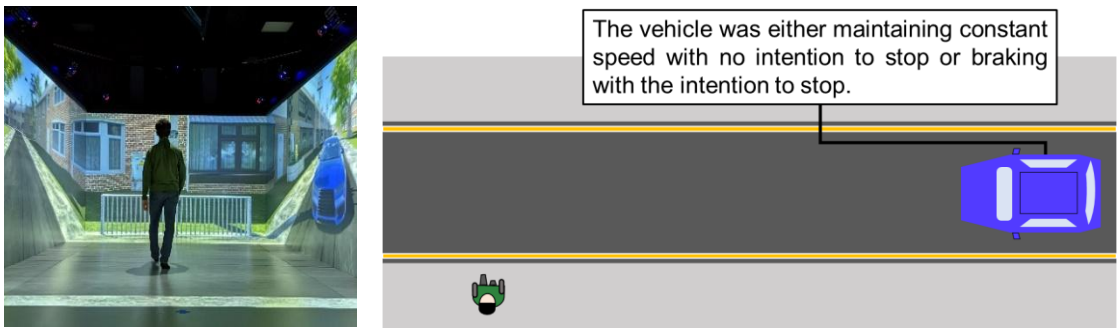

Figure 2.1 - HIKER experimental environment (left) and not-to-scale schematic bird's eye view of the experimental paradigm (right)

2.1.3 Experimental design

The design was adapted from Baker et al. (2009), where observers were shown animations and trajectories of simulated rational agents, that were navigating inside mazes. After observing the agent's movements, the participants had to judge which goal (out of a set of possible goals), was the agent most likely pursuing. That experimental design was modified in accordance with the experiments from Pekkanen et al. (2022), Dey et al. (2019) and Ackermann et al. (2019). More specifically, Pekkanen et al. (2022) designed an experiment in which participants had to decide when it was safe to cross a road as a vehicle approached at either a constant speed or decelerated to a stop. Dey et al. (2019), for instance, measured the willingness of the participants to cross the road in front of an approaching vehicle at predetermined distances or segments. Ackermann et al. (2019), similarly, measured the reaction times of the participants in detecting the decelerating movement of an approaching vehicle. The concept of scenario segmentation was very relevant to the needs of the current work. Following the concepts of the experimental designs of the studies above, the successful cognitive laboratory experimental task by Baker et al. (2009) was expanded to a realistic traffic setting.

The simulated traffic environment included a residential block with a one-lane road that was 4.2 meters wide and an intersection without traffic signals, during daylight hours. A blue sedan vehicle was being controlled by the simulation (its kinematics were predetermined) and was driving in the centre of the road. The vehicle was autonomous (absence of driver and passengers) and at the crossing point there were no markings of a zebra crossing. To focus the study purely on the interpretation of kinematic cues, participants were deliberately not explicitly told that the approaching vehicle was an AV. The purpose of these design choices

was to separate the effects of vehicle kinematics on pedestrian judgments from other factors that might have affected them, such as biases due to driver/passenger behaviour (e.g., making eye contact) or the legal expectations that come with a formal crossing. The aim of the experiment was to investigate how the vehicle's movements would influence the behaviour and judgments of pedestrians, by considering various factors such as the vehicle's driving manoeuvre, the time it takes for the vehicle to reach the pedestrian (initial TTA), and the initial speed of the vehicle. The vehicle approached the pedestrian at three different initial speeds (25 km/h, 40 km/h and 55 km/h) and with two different initial TTAs (3 s and 6 s). Three distinct types of driving manoeuvres were evaluated (illustrated in Figure 2.2):

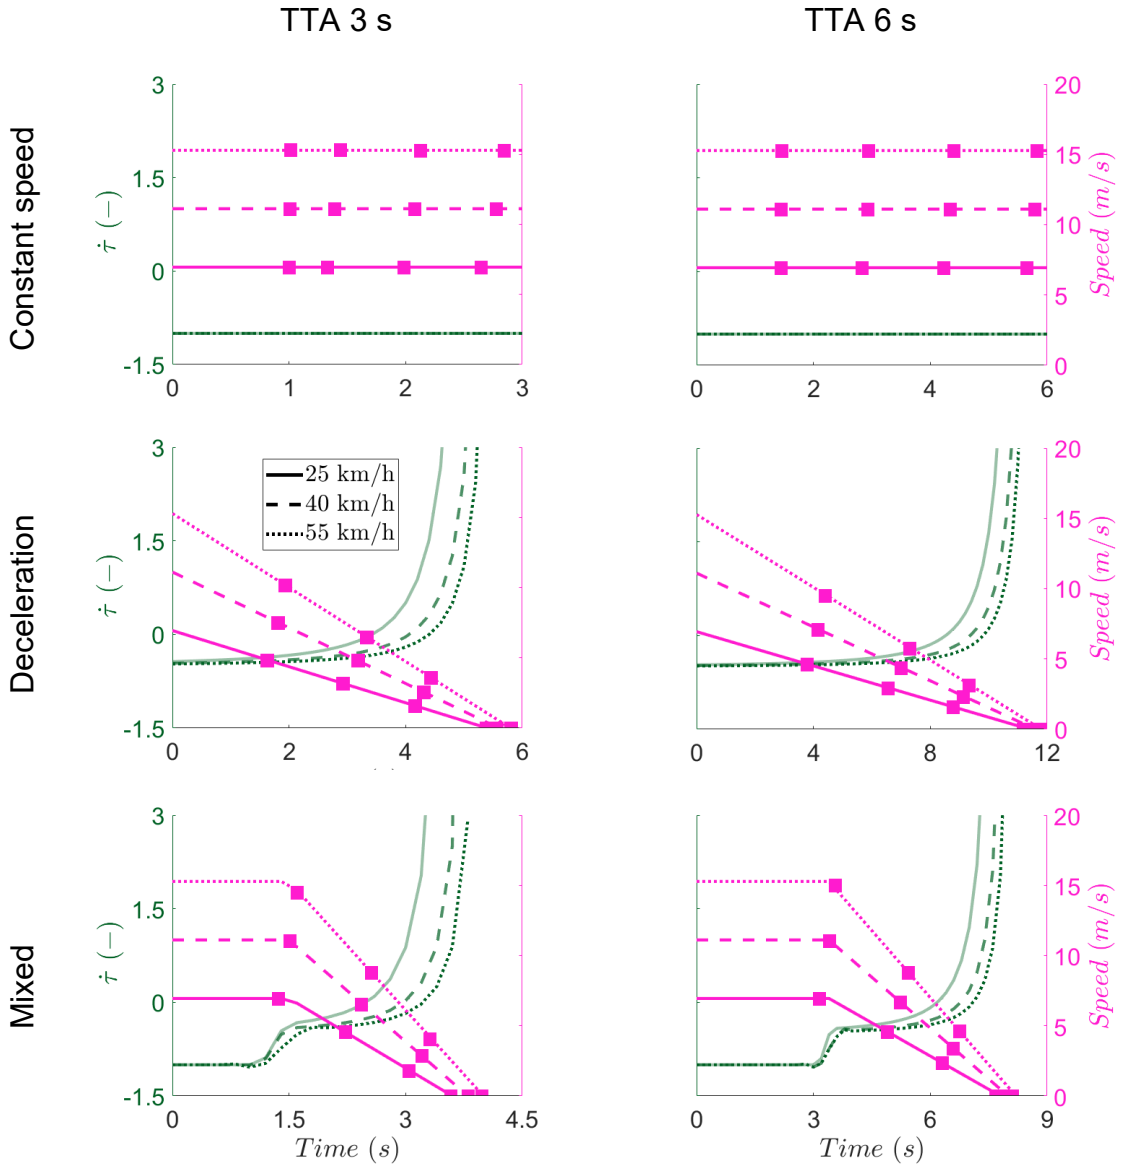

Figure 2.2 - Vehicle kinematics of all 18 kinematic scenarios. The vehicle's speed profile is denoted using the pink curves and the respective $\dot{\tau}$ time history by the dark green curves

- Constant speed: The vehicle maintained a consistent speed throughout the simulation of the scenario.
- Deceleration: The vehicle decreased its speed at a constant rate until it stopped approximately 2.5 m from the participant.
- Mixed: The vehicle combined an initial phase of constant speed and then a deceleration phase, to explore the impact of the vehicle's behaviour change, during the simulation, on the pedestrian's belief updating. In this scenario, the vehicle kept a steady speed for a certain time (1.5 s for the 3 s initial TTA condition and 3.4 s for the 6 s initial TTA condition) before slowing down and coming to a stop approximately 2.5 m from the participant. The mixed manoeuvre scenarios had a higher deceleration rate than the corresponding deceleration scenarios as the initial TTA and speed were kept the same.

The initial TTAs and initial speeds were chosen based on existing traffic safety and perception literature (more details can be found in Subsection 1.3.1) to investigate pedestrian judgments of vehicle behaviour in conditions similar to studies that have tested pedestrian road crossing responses/decisions. More specifically, the TTA range of 3 to 6 s was recognised as the critical window for pedestrian gap acceptance, in which road crossing decisions could be significantly variable. The use of these TTA values allowed the investigation of behaviour estimation under time pressure (3 s) versus sufficient time for comfortable decision making (6 s). The initial speed values were chosen since they would represent a range of speeds commonly found in urban environments. The speed variation was necessary to investigate the influence of spatio-temporal distance on beliefs and to address the phenomenon where pedestrians' TTA estimations are influenced by speed. For a given TTA, a higher speed would mean that the vehicle is further away, allowing to test whether pedestrians rely on time, distance, or a combination of both when estimating the behaviour of an approaching vehicle. The deceleration rates in the pure deceleration scenarios ranged from approximately 0.61 m/s^2 to 2.63 m/s^2. These rates are comparable to those frequently observed when vehicles yield normally to pedestrians (Carlowitz et al., 2024; Yang et al., 2024). For the mixed manoeuvre scenarios, the deceleration rates during the braking phase were higher, ranging from approximately 1.58 m/s^2 to 5.87 m/s^2. Such rates can be observed when vehicles

slow down more urgently, for example, before intersections on roads with higher speed limits or in near-emergency situations (Bokare and Maurya, 2017; Kudarauskas, 2007). In summary, the experiment consisted of three types of manoeuvres, three initial speeds and two initial TTAs, resulting in a total of 18 kinematic scenarios. The details of these scenarios can be seen in Table 2.1.

Table 2.1 - Details of kinematic scenarios

Manoeuvre	**Initial TTA (s)**	**Initial speed (km/h)**	**Initial distance to pedestrian (m)**	**Deceleration rate (m/s^2)**	**Constant speed phase duration (s)**
Constant speed	3	25	20.83	-	[0, 3]
		40	33.33	-	[0, 3]
		55	45.83	-	[0, 3]
	6	25	41.67	-	[0, 6]
		40	66.67	-	[0, 6]
		55	91.67	-	[0, 6]
Deceleration	3	25	20.83	1.29	-
		40	33.33	1.98	-
		55	45.83	2.63	-
	6	25	41.67	0.61	-
		40	66.67	0.97	-
		55	91.67	1.04	-
Mixed	3	25	20.83	3.15	[0, 1.5]
		40	33.33	4.63	[0, 1.5]
		55	45.83	5.87	[0, 1.5]
	6	25	41.67	1.58	[0, 3.4]
		40	66.67	2.42	[0, 3.4]
		55	91.67	3.18	[0, 3.4]

2.1.4 Tasks and procedure

The experiment was consisted of two main tasks which were performed sequentially: a road crossing task (two blocks) followed by a behaviour estimation

task (one block). The behaviour estimation task is the primary focus of the modelling efforts in this thesis. The road crossing task served two purposes: first, to allow participants to familiarise with the virtual environment and immerse them in a realistic road crossing mindset before they undertook the behaviour estimation task and second, to enable potential future comparisons between behaviour estimation judgments and actual crossing decisions, as presented in the work by Tian et al. (2023).

The first block of the road crossing task was preceded by a practice session of 10 trials to familiarise participants with the VR environment and task requirements. For the road crossing task, participants were positioned at a floor marker 57 cm from the kerb. At the start of each trial, their view of the road to the right (from which the vehicle would approach) was initially obscured by a grey virtual plane. A message “Look here. Keep looking” was displayed opposite of the participant (on the other side of the street). After the participants kept looking at the message for a very brief time, the message disappeared. The participants were instructed to turn their head to the right (after the message had disappeared), which caused the grey obstruction plane to become transparent, revealing the approaching vehicle scenario. The participants were then to decide whether to cross, if and when they felt safe. The specific instructions regarding movement were: “If you decide to cross, please walk naturally as you would in everyday life. If you decide not to cross, please remain standing at your initial position and wait until you feel safe to cross.”. The trial concluded once the vehicle had passed the crossing point. Participants then returned to the starting marker if they had crossed or remained there if they had waited. Both blocks of the road crossing task included the same 18 experimental scenarios (illustrated in Figure 2.2), presented in a randomised order within each block. After completing the two road crossing blocks, a short break was provided, before the block of the behaviour estimation task began.

The second task involved assessing a vehicle's behaviour by determining if it was stopping to allow the pedestrian cross or if it was maintaining its speed to pass first. In the behaviour estimation task, the same scenarios as in the road crossing task were reused, but in line with the paradigm in Baker et al. (2009), each of the 18 kinematic scenarios was truncated at 4 different points, creating four different segments of different durations. These segments are illustrated as

square markers on the speed profiles of the 18 kinematic scenarios in Figure 2.2, indicating the points where each scenario was truncated. The shortest of these segments showed just a short time at the start of the scenario in question, giving very limited information based on which the participant could judge the vehicle's behaviour, while the longest segments showed the entire approach of the vehicle.

In the deceleration and mixed scenarios, the first, shortest segment showed little to no stopping evidence, the second and third displayed increasing stopping evidence, and the fourth clearly indicated stopping behaviour. As can be seen in Figure 2.2, the visual cue for collision judgment, $\dot{\tau}$, increases exponentially as the vehicle approaches the pedestrian's position. Therefore, to achieve the intended progression of increasing deceleration cues in longer segments, a logarithmic distance division method was applied, given by:

$$D_i = a^{5-i}; a = \sqrt[5]{D_{int}}, i = 1,2,3 \; and \; D_4 = 2.5 \; m \tag{2.1}$$

Where D_i refers to the distance between the approaching vehicle and the pedestrian at the end of the ith segment, a is the logarithmic base based on the initial distance of the approaching vehicle, D_{int} the distance at the end of the 4th segment of all traffic scenarios equals 2.5 m, i.e., the final stopping distance from pedestrians.

For the constant speed scenarios, the segments were created by evenly dividing the vehicle's approach path into four temporal or spatial portions, as the visual cue for collision judgment, $\dot{\tau}$, remains constant ($\dot{\tau} = -1$). The divisions were calculated as follows:

$$D_i = D_{int} - bi; b = \frac{D_{int} - D_4}{4}, i = 1,2,3 \; and \; D_4 = 2.5 \; m \tag{2.2}$$

Where D_i refers to the distance between the approaching vehicle and the pedestrian at the end of the ith segment, b is the linear base based on the initial distance of the approaching vehicle, D_{int} the distance at the end of the 4th segment of all traffic scenarios equals 2.5 m, i.e., the final stopping distance from pedestrians.

In instances where the calculated duration for the first segment was deemed too short (possibly providing very subtle stimuli or trials being incredibly short for the participant to comprehend), its duration was fixed at a minimum of 1 s. The

segmentation process described above resulted in $18\ scenarios\ x\ 4\ segments/scenario = 72\ total\ segments$.

Similarly to the road crossing task, the behaviour estimation task was also preceded by a practice session of 10 trials using a subset of segments. The trial initiation procedure was identical to the road crossing task: the view to the right was initially obstructed, the “Look here. Keep looking” message was shown on the opposite side of the street, which participants had to look for a brief time, and then participants turned their head right to observe the presented vehicle segment. No road crossing was required, meaning that participants had to remain standing at the floor marker that was used as the starting position in the road crossing task. After each of the 72 segments/trials (presented in a randomised order, ensuring each participant viewed each segment once) was shown, the virtual environment was obscured. Participants were then presented with two questions on the display screen opposite of where they were standing (see Appendix C for question format):

1) “Was the vehicle stopping for you or was it maintaining its speed and passing you?”.
2) “How confident are you in your previous answer? Please rate your confidence level on a scale from 1 to 9.” – Likert scale from 1 to 9 (1 = not confident at all, 5 = somewhat confident, 9 = totally confident).

Upon completion of all experimental tasks, participants were asked to fill out a post-experiment questionnaire. This questionnaire collected demographic information, including age, gender and driving experience.

2.2 From participant judgement to belief probability

To explore the influence of the kinematics variables that were controlled to create the 18 scenarios of the experiment on pedestrians’ beliefs, the participants’ judgements from the behaviour estimation task were analysed. Specifically, a pedestrian’s belief was derived from the combination of the binary choice between the ‘stopping’ or ‘passing’ vehicle behaviour and the confidence rating regarding the first answer. Some examples are provided for further clarification:

- If a participant answered that the vehicle was ‘stopping’ and rated their confidence as 9, that would be translated as 100% belief that the vehicle was stopping.
- If a participant answered the vehicle was ‘passing’ and rated their confidence as 9, that would be translated as 0% belief that the vehicle was stopping (or equally as 100% belief that the vehicle was not stopping).
- If a participant answered that the vehicle was ‘stopping’ and rated their confidence as 1, that would be translated as 50% belief that the vehicle was stopping. In this case only, an answer that the vehicle was ‘passing’ and rated confidence 1, would provide the exact same belief (i.e., 50% belief that the vehicle was stopping, which also means 50% belief that the vehicle was not stopping, indicating maximum uncertainty).
- If a participant answered that the vehicle was ‘stopping’ and rated their confidence as 5, that would be translated as 75% belief that the vehicle was stopping.

After extracting all participants’ beliefs, a 4-way factorial ANOVA (Bao, n.d.) was conducted on these beliefs. The within-subject factors (main effects) investigated were:

- The vehicle manoeuvre (3 levels – Constant speed, Deceleration, Mixed)
- The initial speed (3 levels – 25 km/h, 40 km/h, 55 km/h)
- The initial TTA (2 levels – 3 s, 6 s)
- The segment (4 levels – Segment 1, Segment 2, Segment 3, Segment 4, corresponding to the increasing duration of the vehicle’s approach presented)

To account for between-individual differences, participant ID was included as a random effect in the model. The 4-way ANOVA described above was performed in MATLAB R2022b, utilising the “anovan” function (MATLAB, 2022). This allowed to identify significant main effects of the four factors mentioned above, as well as any significant interactions between them, on pedestrians’ beliefs.

2.3 Model definition

Consider the scenario illustrated in Figure 2.3, where a pedestrian is at the pavement observing an approaching vehicle. An important cognitive task for the pedestrian is to infer the vehicle driver's or AV's intention, specifically whether the vehicle will stop to allow crossing or continue without stopping. The computational models in this thesis were built on a distinction between an approaching vehicle's overarching goal and its immediate behaviour. In the experimental scenarios, the vehicle's goal was assumed to be known to the pedestrian: to continue driving along the road. The uncertainty for the pedestrian though lied in the behaviour that the vehicle would exhibit to achieve this goal. This thesis defined two mutually exclusive behaviours: stopping (i.e., decelerating with the intention to yield) and not stopping (i.e., maintaining speed with the intention to pass). The cognitive process by which a pedestrian infers which of these two behaviours the vehicle is exhibiting was termed behaviour estimation.

Figure 2.3 - Fundamental scenario of a pedestrian at the pavement observing an approaching vehicle

The computational model suggested in this chapter is an adaptation of the Bayesian inverse planning framework proposed by Baker et al. (2009), drawing specifically from their "Model 1". This model assumes that an agent pursues a single, unchanging goal g (or in the current context, behaviour b) throughout an observation sequence. Although this assumption might be limiting, as in reality a driver could change their mind, Baker et al. (2009) applied this model to experiments with both unchanging and changing goals during a trial. Similarly, in this thesis the adapted model was applied to scenarios of both unchanging vehicle behaviour (during constant speed and deceleration manoeuvres) and changing vehicle behaviour (during mixed manoeuvres). A set of two mutually exclusive behaviours that the approaching vehicle can exhibit, $B = \{b_s, b_{ns}\}$, was

defined, representing the vehicle's possible behaviours of 'stopping' or 'not stopping', respectively.

This chapter explores models where the pedestrian observes a kinematic state of the approaching vehicle at discrete time steps t. This observed state, s_t, can be either the vehicle's deceleration rate (d, in m/s²) or the time derivative of the relative rate of optical expansion ($\dot{\tau}$, dimensionless). As mentioned in Subsection 1.5.1, $\dot{\tau}$ essentially gives information of how quickly the time-to-arrival (TTA or τ, in s; the time it takes for an object to reach a point) is changing and researchers have suggested that braking is controlled by it. The calculations of the deceleration rate of the vehicle and $\dot{\tau}$ are presented below. To calculate the deceleration rate during a simulation time step, the following formula was used, assuming that the deceleration rate is constant between the two time points in question:

$$d = \frac{v_f - v_i}{\Delta t} \; or \; d = \frac{v_i^2 - v_f^2}{2S} \tag{2.3}$$

Where, v_f and v_i is the final and initial speed (m/s) of the simulation time step, respectively, Δt is the total time of the time step (s) and S is the total distance travelled during that time step (m).

$$\dot{\tau} = \frac{dTTA}{dt} \; and \; TTA(t) = \frac{D(t)}{v(t)} \tag{2.4}$$

Where, $v(t)$ is the current vehicle speed (m/s) and $D(t)$ is the current distance (m) between the pedestrian and the vehicle. While τ (tau) in ecological psychology is an optical variable defined by the ratio of visual angle to its rate of change ($\tau = \theta/\dot{\theta}$) – a cue pedestrians use to estimate TTA without calculating distance or speed directly (Tian et al., 2020; Lee, 1976; Lobo et al., 2018) – this definition comes with the assumption that it can provide TTA estimations only for small visual angles and approaching objects that are travelling with constant speed. For computational simplicity and to avoid the previous assumption since the simulated vehicle also performed decelerating motion, the model's calculation of τ was set equal to the true, objectively calculated instantaneous TTA. This implementation kept the theoretical meaning of the variable as the current time remaining while avoiding the complex mechanisms of real-world visual angle perception.

The model's aim is to estimate the posterior probability (belief) about the vehicle exhibiting a particular behaviour $b \in B$, given a sequence of T observations, $s_{1:T} = (s_1, s_2, s_3, \dots, s_T)$. This inference is achieved using Bayes' rule. While Baker et al. (2009) include an "Environment" term (w) in their general formulation: $P(g|s_{1:T}, w) \propto P(s_{1:T}|g, w)P(g|w)$ (Equation 5 in their Supplementary material), in this specific pedestrian-vehicle scenario, the broader environmental context is considered constant across the observation period relevant to a single crossing decision. Thus, this term is omitted in the current adapted formulation for clarity, without affecting the core inference process regarding the vehicle's behaviour. The derivation of the suggested model is presented below. The belief at time t is the posterior probability of behaviour b given all observations up to that point:

$$P(b) = P(b|s_{1:t}) \tag{2.5}$$

The term $s_{1:t}$ represents the sequence of all observations up to time t. The current observation s_t can be separated from the past observations $s_{1:t-1}$. Applying Bayes' rule gives the following:

$$\begin{aligned} P(b|s_{1:t}) &= \frac{P(s_t|b, s_{1:t-1})\, P(b, s_{1:t-1})}{P(s_{1:t})} \\ &= \frac{P(s_t|b, s_{1:t-1})\, P(b|s_{1:t-1})\, P(s_{1:t-1})}{P(s_{1:t})} \end{aligned} \tag{2.6}$$

Similarly to the numerator, the observations in the denominator can be separated to current and past ones, so that $P(s_{1:t}) = P(s_t, s_{1:t-1})$. The definition of conditional probability is applied on the denominator.

$$P(b|s_{1:t}) = \frac{P(s_t|b, s_{1:t-1})\, P(b|s_{1:t-1})\, P(s_{1:t-1})}{P(s_t|s_{1:t-1})\, P(s_{1:t-1})} \tag{2.7}$$

Assuming that $P(s_{1:t-1}) \neq 0$ (meaning the sequence of past observations has a non-zero probability of occurring), the term $P(s_{1:t-1})$ can be cancelled.

$$P(b|s_{1:t}) = \frac{P(s_t|b, s_{1:t-1})\, P(b|s_{1:t-1})}{P(s_t|s_{1:t-1})} \tag{2.8}$$

The denominator $P(s_t|s_{1:t-1})$ is the marginal likelihood of observing s_t given all the past observations $s_{1:t-1}$ and is a normalising constant (evidence) so that the posterior probabilities for the possible behaviours $\{b_s, b_{ns}\}$ sum to 1 (discrete case of marginal likelihood):

$$P(s_t|s_{1:t-1}) = \sum_{b'} P(s_t|b', s_{1:t-1})\, P(b'|s_{1:t-1}) \tag{2.9}$$

The probability $P(s_t|b, s_{1:t-1})$ is usually referred to as the measurement model. A Markov assumption is made, according to which the past observations $s_{1:t-1}$ do not provide additional information for predicting the probability of observing s_t given that the vehicle exhibits behaviour b. Due to that assumption and since each of the two possible behaviours is linked with specific observation values, the measurement model can be simplified as shown below:

$$P(s_t|b, s_{1:t-1}) = P(s_t|b) \tag{2.10}$$

Finally, substituting (2.9) and (2.10) back to (2.8) provides the most common form of recursive Bayesian estimation, which is also the formulation of the suggested model.

$$P(b|s_{1:t}) = \frac{P(s_t|b)\, P(b|s_{1:t-1})}{\sum_{b'} P(s_t|b')\, P(b'|s_{1:t-1})} \tag{2.11}$$

Where, $P(b|s_{1:t})$ is the posterior probability (belief) that the vehicle currently exhibits behaviour b given the observations until time t ($s_{1:t}$), $P(s_t|b)$ is the likelihood of the current observation s_t, given the current vehicle behaviour b, $P(b_t|s_{1:t-1})$ is the prior probability that the vehicle exhibited behaviour b given the observations until time $t-1$ ($s_{1:t-1}$) and $\sum_{b'} P(s_t|b')\, P(b'|s_{1:t-1})$ is the evidence or marginal likelihood. For the formulation of the model a classical Bayesian approach was adopted, where the initial prior probability for the two mutually exclusive behaviours is set to be uniform at $t = 0$: $P(b_{\mathrm{s}}) = P(b_{\mathrm{ns}}) = 0.5$. This assumption means that before a pedestrian receives any sensory evidence from the vehicle's kinematics, they are in a state of maximum uncertainty. The implication of this choice is that any subsequent shift in the belief probability is driven only by the accumulation of new kinematic evidence. However, this assumption represents a simplification of real-world cognition. A pedestrian's belief is likely influenced by prior experience, both in the real world (e.g., learned traffic norms at an unmarked crossing) and possibly within the context of the experiment itself (e.g., exposure to previous scenarios). Such experience would establish a learned prior bias that moves the initial belief away from the neutral belief. This limitation is addressed in the next chapter by incorporating a prior bias parameter.

The likelihood term $P(s_t|b)$ quantifies how well the observed kinematics match what is expected for a given behaviour. This probabilistic mapping from an underlying vehicle behaviour b to an observed kinematic state s arises from two main sources:

1) Perceptual Noise: The pedestrian's estimation of the vehicle's deceleration rate d or $\dot{\tau}$ value may not be perfectly accurate. This was modelled as normally distributed noise with a mean of zero and standard deviation (or noise intensity) N_p (one of the free parameters of the model).
2) Behavioural Variability: Vehicles may not exhibit a given behaviour (e.g., stopping) with an identical kinematic profile every time. This differs from the primary source of probability in Baker et al.'s (2009) models, where environmental states are generally known, and agents probabilistically choose actions to achieve goals based on rational planning. In the proposed model, the underlying behaviour b (analogous to Baker et al.'s (2009) goal g) is fixed for the sequence, and the observed s states are noisy manifestations of this fixed behaviour. This type of uncertainty was assigned the standard deviation Ω_b.

To capture this combined uncertainty, it was assumed that the likelihood of observing a particular kinematic state s given a behaviour b, $P(s_t|b)$, follows a normal (Gaussian) probability distribution, $\mathcal{N}(s_t; \mu_b, \sigma_b^2)$. The mean μ_b represents the most typical kinematic value the pedestrian expects for a given behaviour b. Figure 2.4 illustrates examples of these likelihood distributions for deceleration rates and $\dot{\tau}$, showing, for instance, that an observed deceleration rate of 0 m/s^2 or $\dot{\tau} = -1$ might be most probable under a 'not stopping' behaviour, while a deceleration of 4 m/s^2 or $\dot{\tau} = -0.5$ might be most probable under a 'stopping' behaviour. Specifically, the explanations behind these $\dot{\tau}$ values have been provided in Subsection 1.5.1. The standard deviation σ_b of this distribution represents the pedestrians' overall uncertainty for a given behaviour b. The overall uncertainty is decomposed to the two sources described above as:

$$\sigma_b = \sqrt{N_p^2 + \Omega_b^2}$$

These mean and standard deviation parameters define what kinematic values are considered typical for 'stopping' versus 'not stopping' behaviour and how much variability is expected around these typical values. In this study, these

parameters (μ_b and σ_b for both 'stopping' and 'not stopping' behaviours, and for both deceleration rate d and $\dot{\tau}$-based alternatives) were treated as free parameters. All the free parameters' (μ_s, μ_{ns}, σ_s, σ_{ns} and N_p) values were estimated by fitting the overall model to the experimental data collected from the participants (details of the fitting procedure are described in the next section.

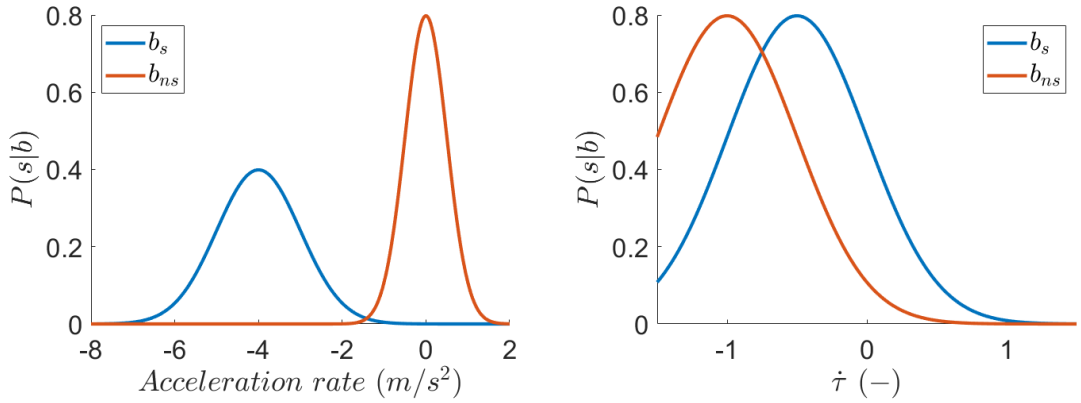

Figure 2.4 - Examples of the $P(s|b)$ probabilities of observing a certain acceleration rate or $\dot{\tau}$ observation given that the vehicle is exhibiting behaviour b. The likelihood was assumed to be following a normal distribution for both possible behaviours

In the framework by Baker et al. (2009), the observed agents select a sequence of actions (e.g., 'move North') probabilistically to achieve a goal state (e.g., 'reach location X'). Therefore, their model involves an inverse planning step by the observer to infer agent goals from observed actions. In the current pedestrian-vehicle context, the vehicle behaviours ('stopping' or 'not stopping') are analogous to these underlying goals. The observable kinematic states (s), are direct, albeit noisy, manifestations of these behaviours. Thus, the model focuses on inferring the behaviour directly from these observed kinematic states, without an equivalent intermediate layer of planning discrete actions as seen in Baker et al.'s (2009) original context.

2.4 Model fitting

To determine the optimal values for the free parameters of the behaviour estimation model, which is based on recursive Bayesian estimation, the experimental data that were gathered from the designed experiment, were used. The model alternatives have multiple parameters. As outlined in the previous

section, the model incorporates several free parameters: the means of the likelihood distributions for the ‘stopping’ and ‘not stopping’ vehicle behaviours (μ_s and μ_{ns} respectively), the standard deviations reflecting the process variability for these behaviours (σ_s and σ_{ns}) and a factor for the observation noise (N_p).

It is noteworthy that μ_{ns} could be fixed on specific values based on physical assumptions for both model alternatives, however in the current fitting procedure, all five listed parameters were treated as free, allowing more flexibility.

1) $\mu_{ns} = 0$ for the d-based alternative since $d = 0\, m/s^2$ if the vehicle is travelling with constant speed
2) $\mu_{ns} = -1$ for the $\dot{\tau}$-based alternative (as also illustrated in Figure 2.2).

For simplicity, the experimental data (participants’ ratings regarding the approaching vehicle’s behaviour) were averaged across all subjects, as done by Baker et al. (2009), and the fit that described the averaged data best, would be selected as the best-fitting parameter combination. However, with averaging the data across all participants, it was assumed that all participants are the same, i.e., can be described by a single model parameterisation. The model output is the estimated posterior belief (belief probabilities of the most-likely current behaviour of the vehicle). Since the behaviour estimation model has more than one free parameter, the multi-parameter model fitting procedure was adopted. A grid search was performed, i.e., all combinations of the ranges of values for each parameter, as shown in Table 2.2, were tested. Since the inspiration was drawn from Baker et al. (2009) in designing the experiment and formulating the model, the model fitting approach was also adopted directly from the methodology used in their cognitive science study. More specifically they compared their models using a bootstrap cross-validated (BSCV) correlational analysis, a non-parametric technique that fit models on random training subsets and tested them on the complementary data. That analysis included a grid search that tested a number of parameter values, ultimately calculating the average correlation between the participants’ data with the testing datasets to assess goodness-of-fit of the models. Additionally, since the suggested model was quite simple, including only five free parameters, a grid model fitting approach allowed an exhaustive search of the parameter ranges.

Table 2.2 - Model's parameter ranges

Parameters	d-based alternative	$\dot{\tau}$-based alternative
μ_s	{-6, -5.5, -5, …, 0}	{-1, -0.75, -0.5, …, 1}
μ_{ns}	{-6, -5.5, -5, …, 0}	{-1, -0.75, -0.5, …, 1}
σ_s	{0.1, 0.4, 0.7, …, 1.3}	{0.1, 0.4, 0.7, …, 1.3}
σ_{ns}	{0.1, 0.4, 0.7, …, 1.3}	{0.1, 0.4, 0.7, …, 1.3}
N_p	{0, 0.5, 1, …, 5}	{0, 0.5, 1, …, 5}

To express participants' ratings, which were originally on a 1 to 9 Likert scale, as probability beliefs, a linear transformation was applied to convert these ratings to a 0-100% scale. To quantify the goodness of fit for each parameter combination, the average correlation coefficient (Spearman's rho) between the model's predictions and the average beliefs of the participants across all 18 kinematic scenarios, was calculated. The Spearman's rank correlation, represented by ρ_S, shows the relationship between two sets of data. Unlike Pearson's correlation that uses the original data values, ρ_S is determined by the ranked order of the data. In the current case, for a dataset X, which represents beliefs predicted by a model, and a dataset Y, which represents the average beliefs of participants, the values in X and Y are converted into their respective ranks, denoted as RX and RY. The Spearman's correlation (ρ_S) is then calculated using these ranks.

$$\rho_S = 1 - \frac{6 \sum d_i^2}{n(n^2 - 1)} \tag{2.12}$$

Where $d_i = RX_i - RY_i$ is the difference between the ranks of corresponding data points, and n is the number of data points.

The choice of Spearman's correlation instead of Pearson's was made, as the former assesses the monotonic relationship between two variables, making it robust even if their connection is not strictly linear and less sensitive to outliers than the latter. This was deemed appropriate as subjective rating scales may not perfectly map to linear changes in model probability outputs. Spearman's correlation captures the following relationship: if a model probability increases, subjective rating also tends to increase (or decrease, if negatively correlated), without assuming a linear transformation. Finally, the parameter combination

yielding the maximum average correlation was selected as the best-fitting parameter combination.

2.5 Results

This section presents the key findings from the experiment, performing a statistical analysis of participants' beliefs across the different experimental conditions. Following, the performance of the proposed computational model alternatives in predicting these subjective ratings is evaluated.

2.5.1 Pedestrian beliefs regarding the vehicle's behaviour

2.5.1.1 Segments and driving manoeuvres

Figure 2.5 illustrates participants' beliefs (P_s, indicating the belief that the vehicle is stopping) as a function of the presented segment and the vehicle's driving manoeuvre. A 4-way factorial ANOVA was conducted as described in Section 2.2, to analyse those beliefs. The full table of this analysis is provided in Appendix A . The analysis revealed a significant main effect of segment on pedestrians' beliefs $\left(F(3,1172) = 24.54,\ p < .001,\ \eta_p^2 = 0.130\right)$, suggesting that as more time passed and the participants observed a vehicle's approach for longer, the more certain (beliefs closer to 100% or 0%) they were of whether the vehicle was stopping or not. Pedestrians' beliefs seemed to be affected by different vehicle driving manoeuvres, indicated by a significant main effect $\left(F(2,1172) = 589.59,\ p < .001,\ \eta_p^2 = 0.782\right)$ of driving manoeuvre. Notably, pedestrians generally believed that the vehicle was stopping, with higher certainty, during Deceleration manoeuvres than during Constant speed or Mixed manoeuvres. At the same time pedestrians generally believed that the vehicle was not stopping, with higher certainty, during Constant speed manoeuvres than during Deceleration or Mixed manoeuvres.

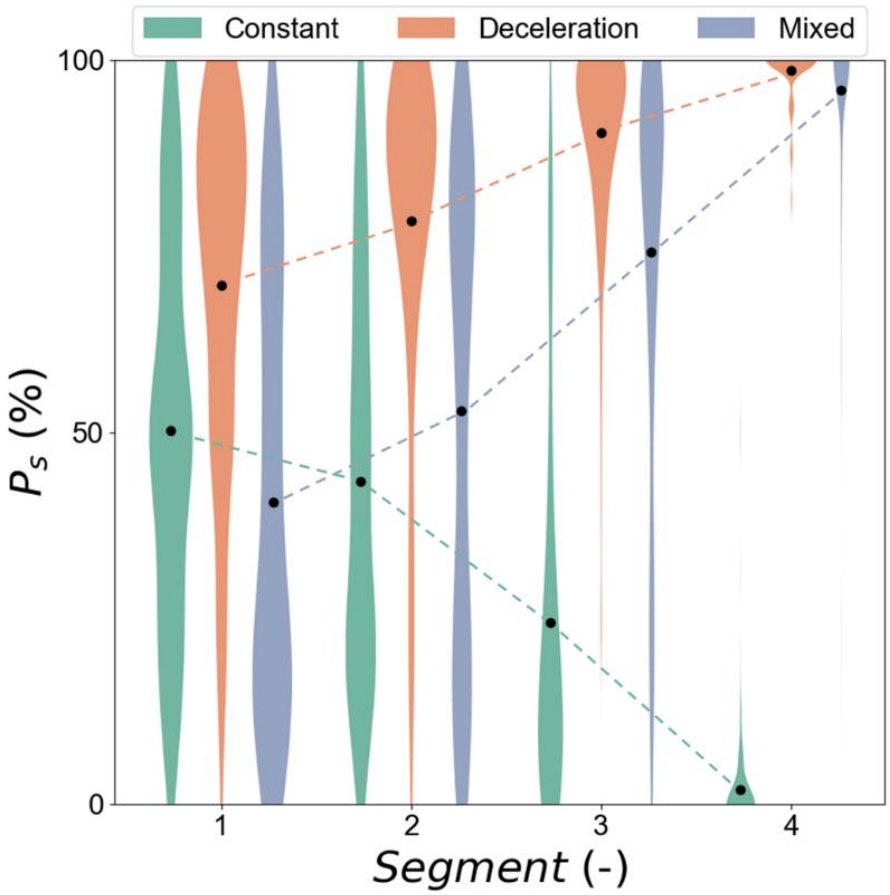

Figure 2.5 - Pedestrians' beliefs (P_s) regarding that the vehicle is stopping as a function of segment and the presented driving manoeuvre. Means are represented by black dots and dashed lines indicate the trends of the average beliefs

For the Constant speed manoeuvres (green violins and mean trendline in Figure 2.5) participants' belief that the vehicle is not stopping strengthened as more of the scenario was revealed over time. Initially (Segment 1), the average belief was around 50%, indicating uncertainty, with considerable variability in participant responses. This variability can be explained by the fact that all different initial speed and initial TTA levels are grouped under the umbrella of a specific driving manoeuvre, in this case the Constant speed one. As the vehicle continued at a constant speed through subsequent segments, this belief shifted even more decisively towards the vehicle not stopping (closer to $P_s = 0\ \%$) and the variability was decreased, as it was more obvious that the vehicle was not stopping in later segments, for all different initial speeds and TTAs. Targeted post-hoc pairwise comparisons were conducted to examine differences between successive time segments (Segment 1 vs. 2, Segment 2 vs. 3, and Segment 3 vs. 4) within the constant speed manoeuvres (Appendix A). The results showed that the average

beliefs of the pedestrians had statistically significant differences for these sequential time segment comparisons.

Conversely, when the vehicle was observed to be decelerating from the start of perceived approach (orange violins and mean trendline in Figure 2.5), pedestrians generally identified the stopping behaviour, and their belief in this (P_s) increased (became more certain) over longer segments. Even in segment 1, the mean belief was already leaning towards the stopping behaviour and the trend continued towards high certainty by segment 4. The higher to lower variability trend can be explained in a similar manner as in the Constant speed manoeuvres. As in the Constant speed manoeuvres, targeted post-hoc analysis showed that the average beliefs of the pedestrians had statistically significant differences for sequential time segment comparisons in the Deceleration manoeuvres (Appendix A), as well.

As shown in Figure 2.5, in Mixed scenarios (blue violins and mean trendline), where the vehicle initially (segment 1) maintained constant speed and then (rest of the segments) decelerated, pedestrians' beliefs evolved in a different way than in the other two driving manoeuvres. In early segments, pedestrians seemed to be, on average, uncertain or were slightly leaning towards the non-stopping behaviour (consistent with the vehicle's initial constant speed phase). As the vehicle began to decelerate in later segments, pedestrians' beliefs shifted towards the vehicle stopping behaviour. However, pedestrians' beliefs in Mixed manoeuvres often were not reaching the same level of certainty seen in the Deceleration manoeuvres before the equivalent final segment. Targeted post-hoc analysis on the Mixed manoeuvres confirmed once again statistically significant differences in average beliefs between the successive segments (Appendix A).

Finally, a significant interaction effect involving the segment and the driving manoeuvre on pedestrians' beliefs regarding the vehicle's behaviour $\left(F(6,1172) = 202.92,\ p < .001,\ \eta_p^2 = 0.668\right)$, was found. A targeted post-hoc analysis was conducted to compare the average pedestrian beliefs between different driving manoeuvres at each of the four segment levels (Appendix A). This analysis showed that the average beliefs of the pedestrians had statistically significant differences between different driving manoeuvres, at the same segment levels, except in the case of the comparison between the 4th segments

of the Deceleration and the Mixed manoeuvres, since at the end of both of these manoeuvres it was quite evident that the vehicle was stopping.

2.5.1.2 Initial TTAs and initial speeds

Figure 2.6 illustrates how pedestrians' beliefs regarding the approaching vehicle's behaviour was influenced by the vehicle's initial TTA (left column panels) and the vehicle's initial speed (right column panels). The 4-way ANOVA (results for all factors and interactions detailed in Appendix A) revealed significant main effects of initial TTA $\left(F(1,1172) = 98.67,\ p < .001,\ \eta_p^2 = 0.158\right)$ and initial speed $\left(F(2,1172) = 132.68,\ p < .001,\ \eta_p^2 = 0.302\right)$. In regard to the initial TTA, as also seen in Figure 2.6, longer initial TTAs generally led pedestrians' beliefs to lean more towards the vehicle engaging in stopping behaviour. Conversely to TTA, higher initial speeds, generally led pedestrians to believe more strongly that the vehicle was not stopping. The initial TTA had a positive effect on the average belief of the pedestrians, regarding the decelerating behaviour of the vehicle. These patterns are consistent for all three driving manoeuvres, as shown in the Figure 2.6. The 4-way ANOVA also identified a significant interaction between the initial TTA and initial speed $\left(F(2,1172) = 3.28,\ p = .0448,\ \eta_p^2 = 0.006\right)$, suggesting their combined influence on the average pedestrian beliefs.

More significant interaction effects involving initial TTA, initial speed, segment and driving manoeuvre were found by the 4-way ANOVA, indicating that the way pedestrians' beliefs were influenced by a single factor also depended on the specifics of the other factors. While Figure 2.5 and Figure 2.6 primarily highlight the overall trends of the pedestrians' beliefs for the four factors mentioned above, these more complex interactions, even though they are not easily visualised and comprehended, are statistically supported by the full ANOVA (Appendix A).

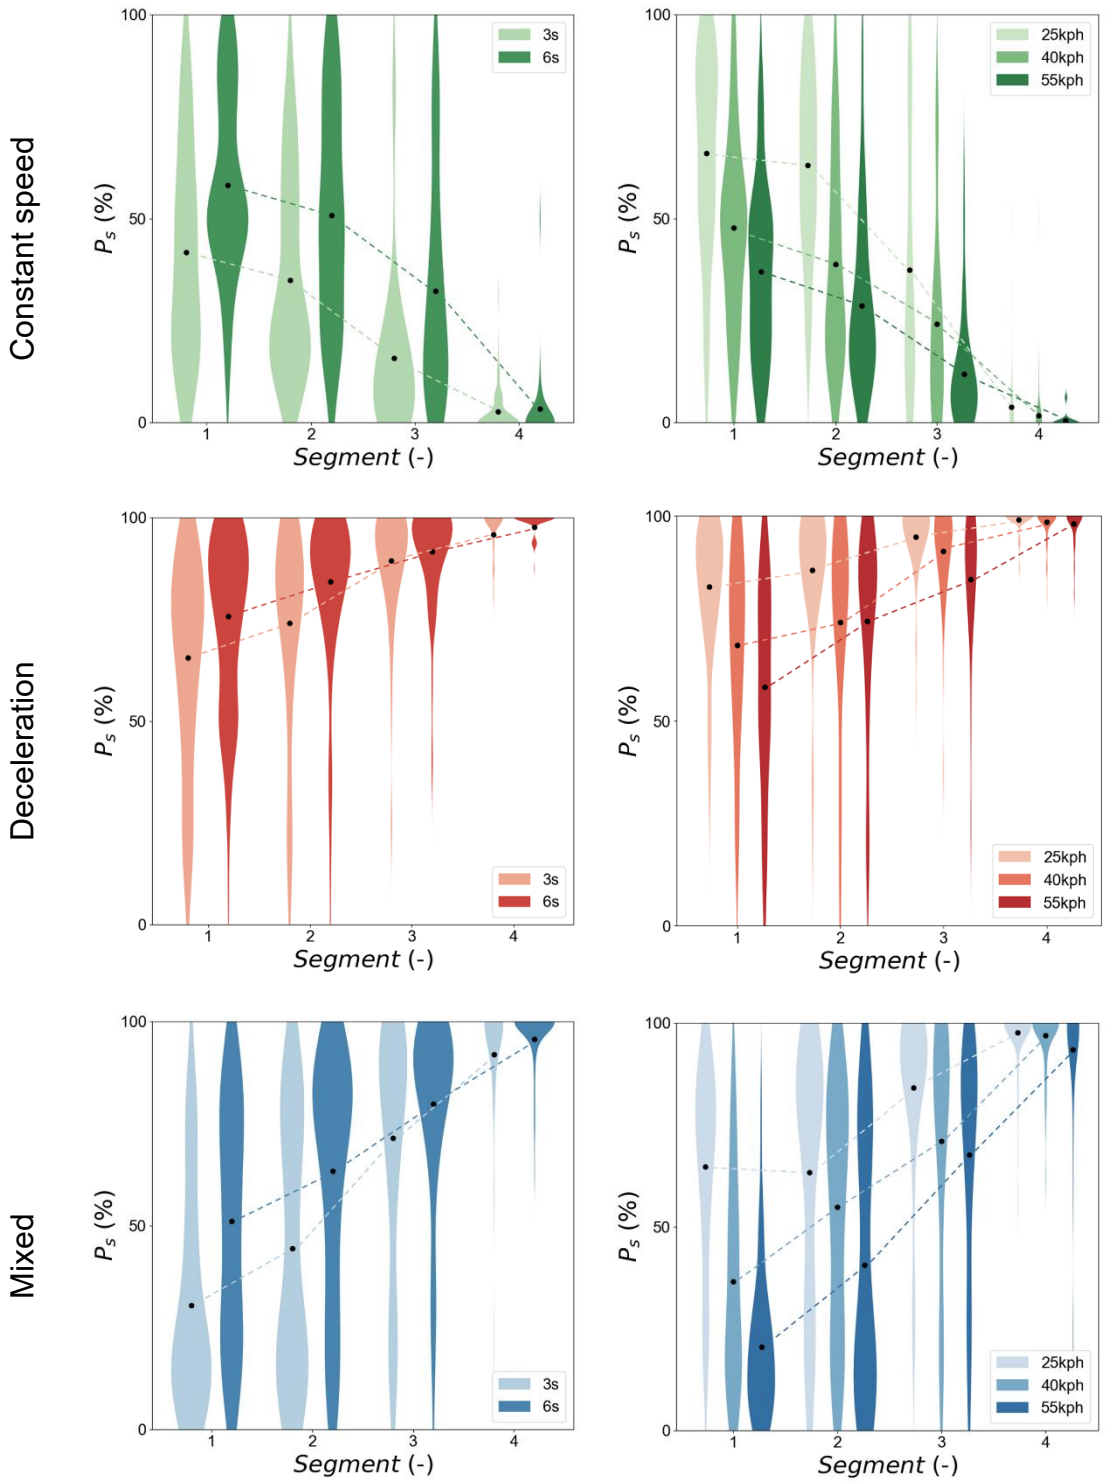

Figure 2.6 - Pedestrians' beliefs regarding the approaching vehicle's behaviour as a function of judgment point, initial TTA (left) and initial speed (right) of the vehicle

2.5.2 Model predictions vs subjective ratings

Spearman's rank correlation (ρ_S) was used to quantify the relationship between model predictions and the average pedestrian beliefs (participants' subjective ratings). With ρ_S, the goal was to assess the strength and direction (positive or negative) of the relationship between the two variables in question

(model predictions and average pedestrian beliefs), or in other words if the model predictions can capture the trends of the average pedestrian beliefs. A $\rho_S = 1$ would indicate that the model makes belief predictions which rank exactly as the participants' beliefs, across all 72 segments. On the contrary, a $\rho_S = 0$ would indicate that the ranks of model predicted beliefs do not covary with the ranks of the participants' beliefs, i.e., there is no monotonic relationship between the two (for example, a rank increase of participant's beliefs would not be followed by a consistent rank increase of model predictions). The model predictions that are presented in this subsection have been obtained by the parameter combinations which yielded the highest Spearman's correlation between the model predictions and the average pedestrian beliefs. More specifically, the best-fitting parameter settings for both model alternatives are presented in Table 2.3.

Table 2.3 - Best-fitting parameter settings

	d-based alternative	**$\dot{\tau}$-based alternative**
μ_s	-1 m/s^2	-0.5
μ_{ns}	0 m/s^2	-1
σ_s	1 m/s^2	1
σ_{ns}	1 m/s^2	1
N_p	1	0

Figure 2.7 shows the performance of the two model alternatives (using the best-fitting parameter values as stated above), by displaying scatter plots of the model predicted beliefs compared to participants' average ratings (translated in belief probabilities). As can be seen from the two scatter plots, even though both model alternatives provided Spearman's correlations (performance metric), which are closer to 1 (i.e., perfect rank order relationship, not necessarily linear as indicated by the identity line and Pearson's correlation) than 0 (i.e., no rank order relationship), between their belief predictions and the participants' subjective ratings, the $\dot{\tau}$-based model alternative achieved a higher rank correlation. That means that both model alternatives were able to at least capture the overall trends of the beliefs of the pedestrians, regarding the approaching vehicle's behaviour, with the $\dot{\tau}$-based model alternative performing better in that aspect. Moreover, $\dot{\tau}$-based model alternative's better prediction performance is

reflected by a smaller root mean square error (RMSE) between its predictions and the average pedestrian beliefs. Qualitatively, the d-based alternative exhibits a lot more poorly predicted datapoints than the $\dot{\tau}$-based alternative. Additionally, the latter's scatter points are more closely gathered to the identity line, suggesting that its predictions are better estimates than the respective ones of the d-based alternative.

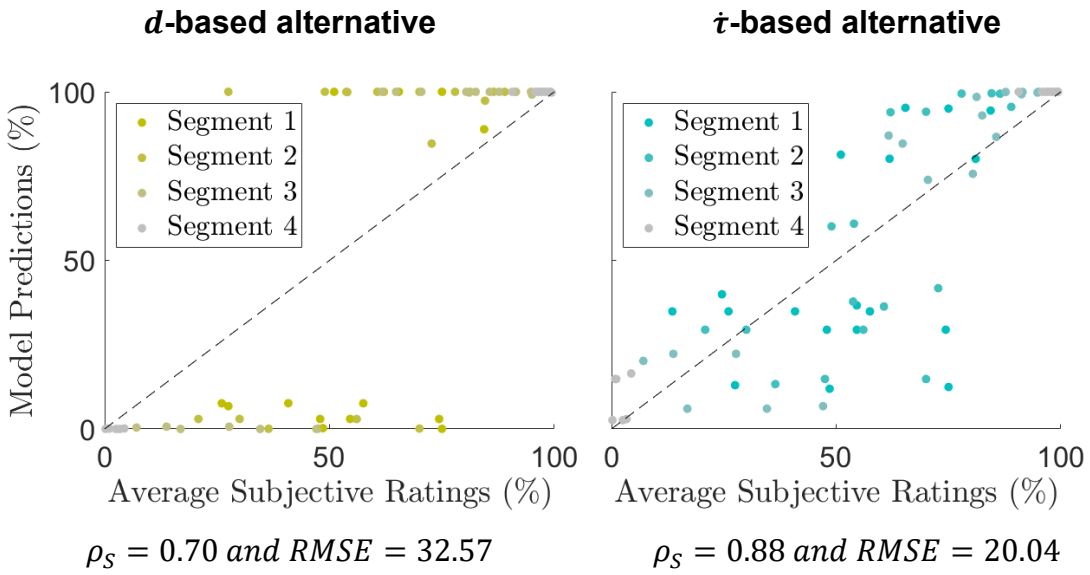

Figure 2.7 - Scatter plots of model predictions using best-fitting parameter settings (y-axes) versus pedestrians' average beliefs (x-axes) for all 18 kinematic scenarios of the approaching vehicle

Following, a qualitative analysis which compares the model alternatives predicted beliefs and the average pedestrians' beliefs, is presented. Through this analysis, the predictive performance of both model alternatives was assessed on each specific kinematic scenario (Figure 2.2 illustrates the kinematics of all 18 scenarios), with each one having a different combination of conditions (driving manoeuvre, initial speed, initial TTA). The presentation below is divided into the scenarios under the three driving manoeuvres (Constant speed – Figure 2.8, Deceleration – Figure 2.9 and Mixed – Figure 2.10). Similarly to Figure 2.7, the d-based and $\dot{\tau}$-based alternative's predictions are shown in yellow and blue coloured curves, respectively. The average pedestrian beliefs are illustrated as standard black standard error of the mean bars at the corresponding segment timings (as described in Subsection 2.1.4).

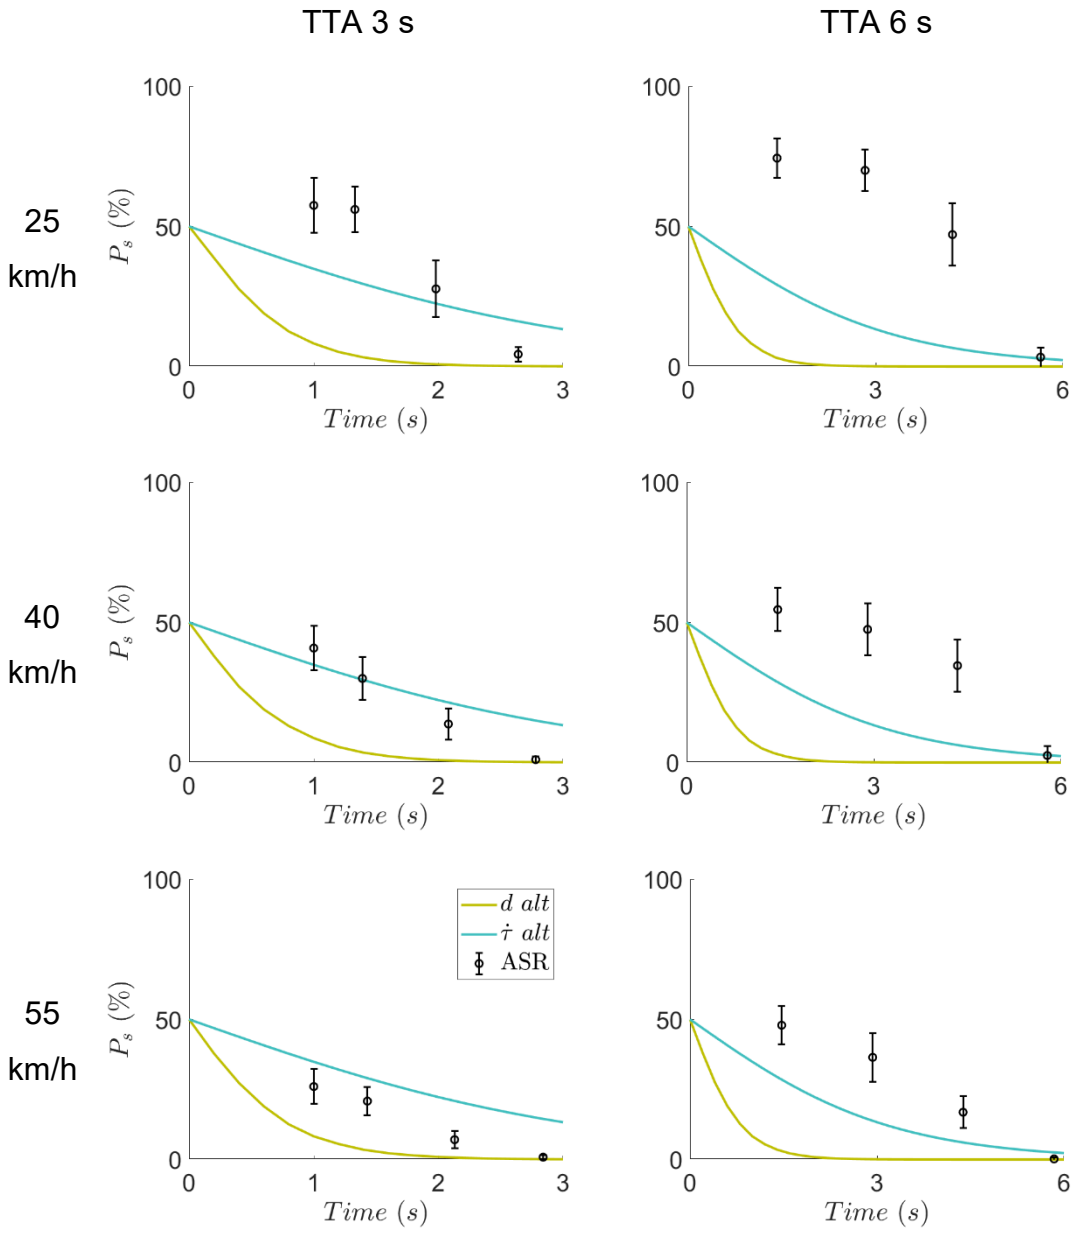

Figure 2.8 - Comparison between model alternatives' predictions and pedestrians' beliefs for the constant speed manoeuvres

As seen in Figure 2.8, in the constant speed manoeuvres, both model alternatives struggled to fully capture the patterns of the average pedestrian beliefs regarding the behaviour of the approaching vehicle. Notably, each model alternative produced a single specific belief curve that is repeated across all six constant speed scenarios. This occurs because the observation state input for both model alternatives is constant in these scenarios ($d = 0\ m/s^2$ and $\dot{\tau} = -1$). Consequently, the models' $P(s|b)$ probabilities remain constant throughout these scenarios, leading to the same belief curve (P_s), regardless of the different kinematic conditions. In contrast, as discussed in Subsection 2.5.1, participants' early judgements in these scenarios were affected by initial speed and TTA. For

instance, with lower initial speeds or longer initial TTAs, participants initially leaned more toward believing the vehicle was stopping. The models' difficulty in capturing these patterns highlights that their current inputs are insufficient for these conditions where these cues are not differentiated. This suggests that pedestrians utilise other information, such as speed and/or spatiotemporal distance, particularly in early segments when d or $\dot{\tau}$ cues are not yet informative. While both alternatives have limitations here, the $\dot{\tau}$-based alternative generally aligns better with the average pedestrian beliefs.

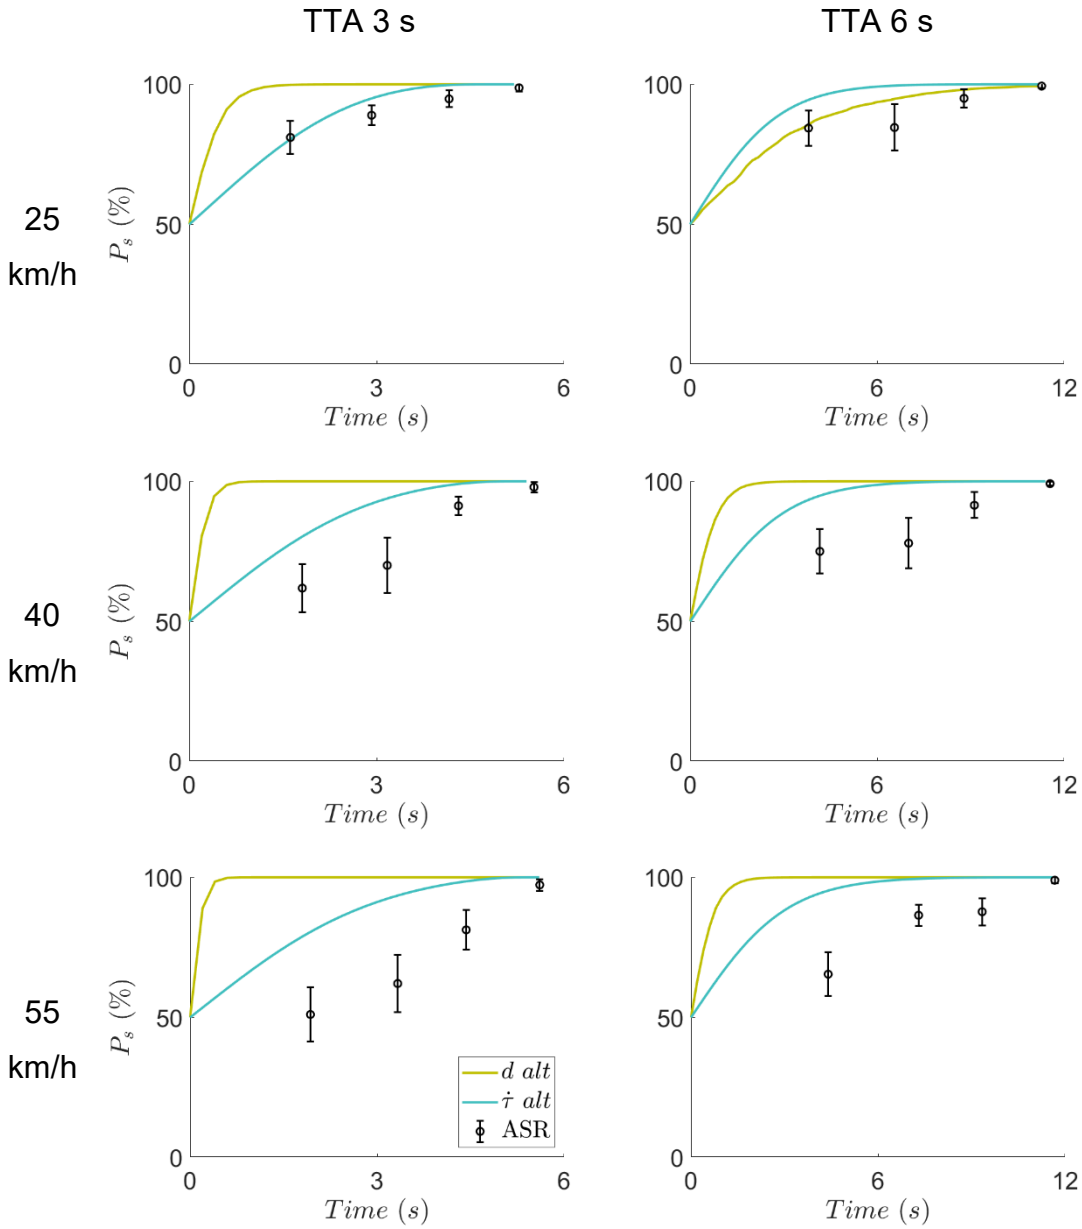

Figure 2.9 - Comparison between model alternatives' predictions and pedestrians' beliefs for the deceleration manoeuvres

In the deceleration manoeuvres, as illustrated in Figure 2.9 and discussed in Subsection 2.5.1, participants' early judgments were, as with the constant speed

scenarios, influenced by the initial speed and TTA. Both model alternatives demonstrated a better ability to capture the overall pattern of increasing belief in the vehicle stopping compared to the constant speed scenarios. As deceleration cues became available, the models captured the human belief patterns more effectively. A tendency for the d-based alternative, to saturate belief at 0% or 100% very quickly was observed; 1) the higher the deceleration amplitude, the faster the belief saturation and 2) this might reflect the sensitivity of the chosen standard deviation parameters of the $P(s|b)$. Later, during the model fitting process it was confirmed (as presented in Figure 2.11) that the optimal values for the σ parameters were found within the defined search ranges (Table 2.2) and were not at the boundaries of these ranges. This suggests that model predictions, with rapid saturation behaviour, is a characteristic of the current model formulation with its best-fitting parameters, rather than a result of a small grid search. The $\dot{\tau}$-based alternative, again, generally showed a better agreement to the average pedestrian beliefs across the different scenarios within this manoeuvre type.

For the mixed manoeuvres (Figure 2.10), where vehicles initially travelled at a constant speed before decelerating, both model alternatives captured the general trends of average pedestrian beliefs, even with this distinct change in vehicle behaviour happening mid-scenario. The $\dot{\tau}$-based alternative once again showed a better performance in capturing the patterns of the average pedestrian beliefs. Notably, the patterns observed in pedestrian beliefs and model predictions during these mixed manoeuvres mirrored those seen in the other two, single-behaviour manoeuvres. Specifically, during the initial constant speed phase of the mixed manoeuvres, the same early belief characteristics and model limitations observed in the Constant speed scenarios were apparent. However, once the vehicle began to decelerate in the later segments of the mixed manoeuvre, the predicted beliefs resembled more to those observed in the pure Deceleration scenarios, with beliefs shifting towards the vehicle stopping. One challenging scenario for both model alternatives was the low initial speed and long initial TTA scenario (25 km/h and 6 s TTA), where average pedestrian beliefs suggest that they believed that the vehicle was stopping, possibly due to its low speed and big distance from them, when it was actually exhibiting a non-stopping behaviour (maintaining constant speed). Both model alternatives, having

knowledge of the deceleration-related kinematics of the vehicle, predicted that pedestrians were extremely accurate at inferring the true behaviour of the vehicle, when in fact they were not, and they were basing their beliefs on other information/expectations; a nuance that both alternatives were not able to capture.

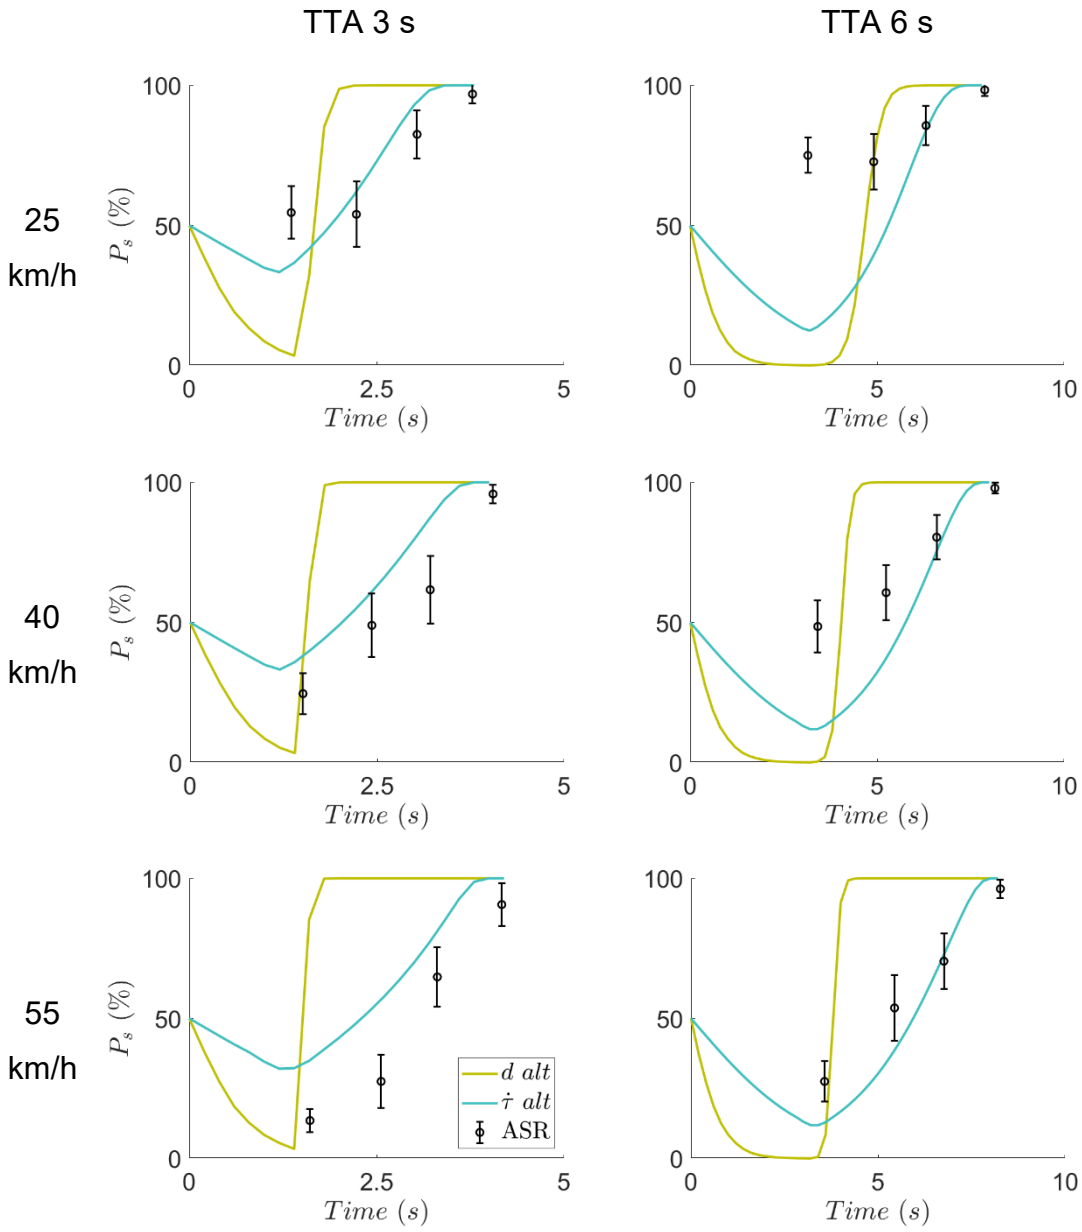

Figure 2.10 - Comparison between model alternatives' predictions and pedestrians' beliefs for the mixed manoeuvres

2.5.3 Parameter analysis

Having established the predictive performance of the model alternatives, this section delves into an investigation of their fitted parameters. The goals of this analysis are threefold: to understand the best-fitting parameter settings

(presented in Table 2.3), to investigate how model fits are affected by variations of parameter settings and to identify potential correlations between parameters, which could indicate issues like parameter redundancy. This analysis is illustrated by pairwise parameter scatterplot matrices for the d-based alternative (Figure 2.11) and the $\dot{\tau}$-based alternative (Figure 2.12).

For the deceleration-rate-based model, the best-fitting parameters were $\mu_s = -1\ m/s^2$, $\mu_{ns} = 0\ m/s^2$, $\sigma_s = 1\ m/s^2$, $\sigma_{ns} = 1\ m/s^2$ and $N_p = 1\ m/s^2$. The mean for the non-stopping behaviour is physically intuitive, directly representing a vehicle maintaining constant speed (i.e., zero acceleration/deceleration). The mean for stopping behaviour reflects a typical, constant deceleration rate, that is within the range of deceleration rates observed in the designed scenarios of the experiment. The equal spread of the likelihood probabilities $P(s|b_s)$ and $P(s|b_{ns})$ is presented in the left panel of Figure 2.13.

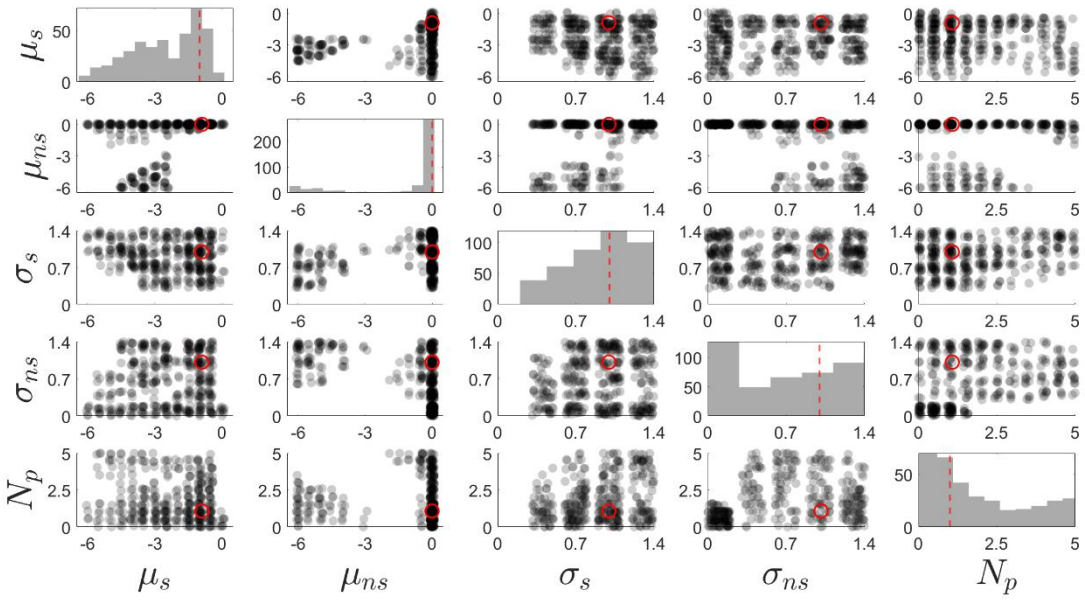

Figure 2.11 - d-based model alternative's pairwise parameter scatterplot matrix. The histograms in the diagonal illustrate the one-dimensional marginal distributions for each parameter, which achieved the highest correlations (above a threshold of 0.65). The scatterplots show the respective two-dimensional marginal distributions for each combination of two parameters, which are represented by black dots (with minor jitter added to improve visibility). The red circles and lines indicate the best-fitting parameter values (the highest achieved correlation was equal to 0.70). The x and y ranges of the scatterplots indicate the full search ranges of the respective parameters

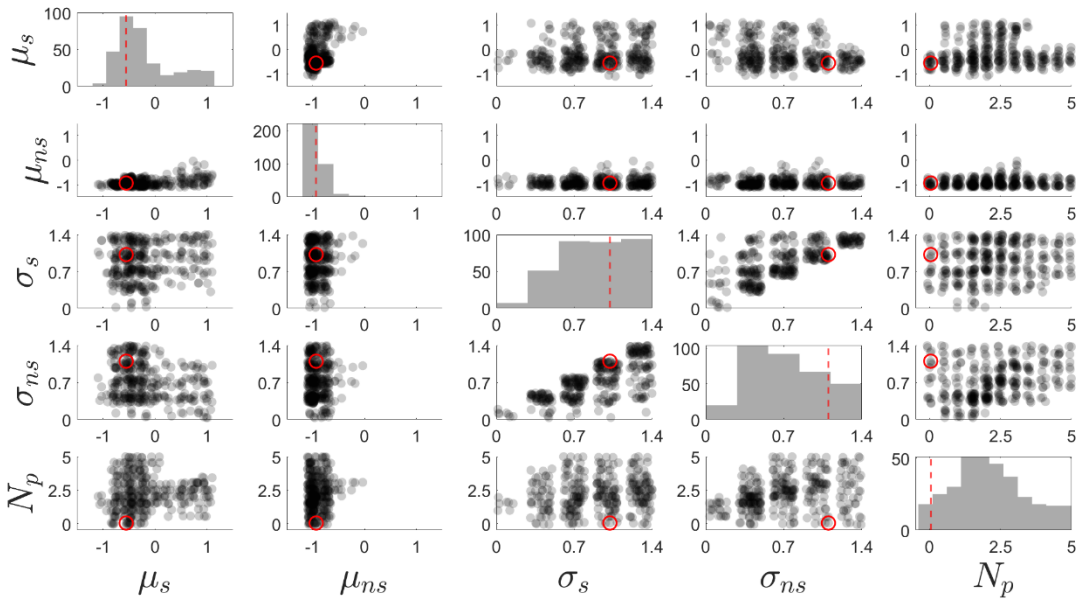

Figure 2.12 - $\dot{\tau}$-based model alternative's pairwise parameter scatterplot matrix. The histograms in the diagonal illustrate the one-dimensional marginal distributions for each parameter, which achieved the highest correlations (above a threshold of 0.8). The scatterplots show the respective two-dimensional marginal distributions for each combination of two parameters, which are represented by black dots (with minor jitter added to improve visibility). The red circles and lines indicate the best-fitting parameter values (the highest achieved correlation was equal to 0.88). The x and y ranges of the scatterplots indicate the full search ranges of the respective parameters

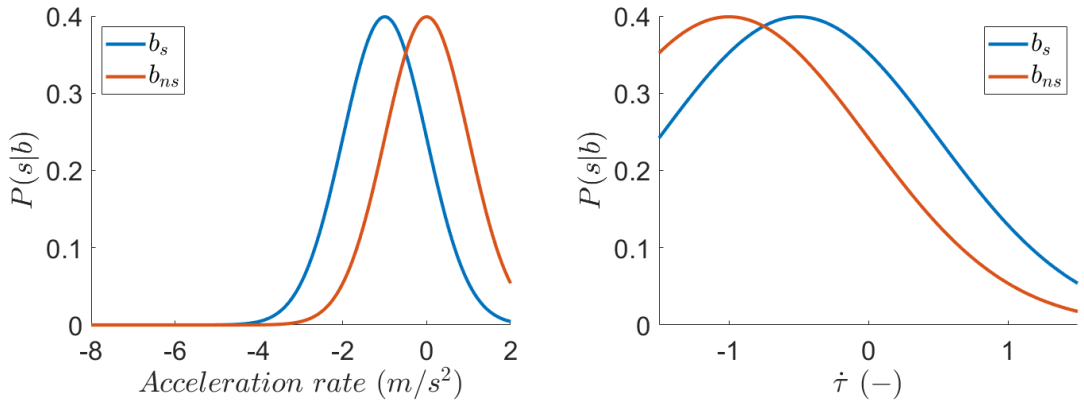

Figure 2.13 - $P(s|b)$ probabilities of observing a certain acceleration rate or $\dot{\tau}$ given that the vehicle is exhibiting behaviour b. The blue (b_s) and orange (b_{ns}) curves refer to the stopping and non-stopping vehicle behaviours, respectively

For instance, with the fitted parameters for the d-based model alternative, the two likelihood distributions intersect at -0.5 m/s^2. This means that any observed acceleration rate lower than this value makes the 'stopping' behaviour more probable than the 'non-stopping' behaviour ($P(s|b_s) > P(s|b_{ns})$), which directs the belief towards the vehicle stopping. Similarly, for the $\dot{\tau}$-based alternative, the cutoff value where the one behaviour becomes more probable than the other is equal to -0.75. Essentially, the μ parameters control the distance between $P(s|b_s)$

and $P(s|b_{ns})$ and the σ parameters control how wide or narrow $P(s|b_s)$ and $P(s|b_{ns})$ are, around their respective means. The combination of all of these indicates a parameter interplay that dictates the model's eventual output. A larger separation between means and smaller standard deviations lead to more certain belief updates, as linking observations to the two possible behaviours is less ambiguous. Conversely, greater overlap between the likelihood distributions results in more uncertainty and slower changes in belief.

To understand the model's sensitivity to its parameters and to identify potential relationships between them, the parameter space was investigated. As illustrated in Figure 2.11, the regions of parameter values with the best fits (highest achieved Spearman's rank correlations) for the d-based alternative, were found around $\mu_s = -1\ m/s^2$ and $\mu_{ns} = 0\ m/s^2$ (where the best fit also resided). As for the standard deviations of the likelihood probabilities there is no clear relationship that can be reported. Finally, no conclusive observations can be made for N_p either, apart from the fact that there was a tendency for lower perceptual noise intensities that yielded higher correlations.

The $\dot{\tau}$-based alternative, which demonstrated better overall performance, yielded the following best-fitting parameters, were $\mu_s = -0.5\ m/s^2$, $\mu_{ns} = -1\ m/s^2$, $\sigma_s = 1\ m/s^2$, $\sigma_{ns} = 1\ m/s^2$ and $N_p = 0\ m/s^2$. Notably, these optimal mean values and the respective high correlation-yielding areas for the likelihood distributions are consistent with established concepts in collision avoidance theory (Lee, 1976) and validate modelling choices of previous pedestrian behaviour modelling studies that have utilised similar $\dot{\tau}$ critical values for modelling pedestrians' crossing decisions (Giles et al., 2019; Markkula et al., 2018; Pekkanen et al., 2022). As mentioned in Subsection 1.5.1, a $\dot{\tau}$ value of -1 signifies no change in the rate of approach, consistent with non-stopping behaviour, while $\dot{\tau}$ values around -0.5 (or more positive) could be interpreted as indicative of significant braking or an intention to yield, as they represent an adequate braking effort from the driver's perspective to stop the vehicle before a pedestrian. Furthermore, an important characteristic of the $\dot{\tau}$ cue is its situation-adaptive nature: during any vehicle yielding manoeuvre with adequate braking, the value of $\dot{\tau}$ is the same (equal to -0.5), providing the same general vehicle behaviour information, regardless of the distance and speed of the vehicle. This contrasts with the deceleration rate, where different distances and speeds would

require different deceleration rate values to stop before the pedestrian's position. Another possible advantage of $\dot{\tau}$ over the deceleration rate could be the fact that when the car is stopping and is in a close proximity to pedestrian, $\dot{\tau}$ increases towards infinity (as shown in Figure 2.2), whereas the deceleration rate is kept constant up until the vehicle completely stops. The abovementioned characteristics of $\dot{\tau}$, offer it an advantage against d as a perceptual cue for inferring vehicle stopping intentions. So, the fact that theoretically grounded mean parameter values were obtained from the model fitting procedure, supports the psychological validity of the $\dot{\tau}$-based alternative.

An examination of the pairwise scatterplots in Figure 2.12 helps to identify potential parameter redundancies. This occurs when different combinations of parameter values can produce very similar model performance, making it difficult to uniquely identify the contribution of each individual parameter. More specifically, the best fits (highest achieved Spearman's rank correlations) were found in the parameter regions of $\mu_s \geq -0.5$ and $\mu_{ns} = -1$, values which are theoretically grounded as stated before. As for the σ parameters, there seems to be a relationship between them; high correlations can be obtained by combinations of lower σ_s and higher σ_{ns} (or vice versa), which indicates a likely redundancy in the parametrisation (for example only one σ parameter for both behaviours being sufficient). The fact that in the $\dot{\tau}$-based alternative the best-fitted noise intensity parameter is equal to zero ($N_p = 0$), might suggest this model alternative can account for the uncertainty in the participants' judgements using only the behavioural variability component of the likelihood function, in comparison to the d-based alternative ($N_p = 1$), which seemed to require the uncertainty of perceptual noise to explain the data a bit better.

2.6 Discussion

The work presented in this chapter aimed to investigate how different initial vehicle speeds, initial Times-To-Arrival (TTAs) and distinct driving manoeuvres influence pedestrians' beliefs regarding an approaching vehicle's behaviour. Then, the ability of a Bayesian observer model, adapted from Baker et al. (2009), to capture the belief updating exhibited by humans in these vehicle approach

scenarios, was explored. This discussion will begin by interpreting the key behavioural observations regarding pedestrian beliefs, particularly the impact of kinematic variables and the temporal dynamics of belief updating. It will then evaluate the performance of the computational model (with two proposed alternatives), addressing both their successes in capturing pedestrians' beliefs and their limitations.

2.6.1 Behavioural observations

A key contribution of the work presented in this chapter is the detailed temporal analysis of pedestrian belief updating regarding vehicle behaviour. By adapting the segmented stimulus presentation approach, inspired by Baker et al. (2009), to a road crossing context, it was possible to investigate how beliefs evolved as more kinematic information became available. This contrasts with past research that, while addressing related topics like gap acceptance or crossing decisions, e.g., Beggiato et al. (2017); Dey et al. (2019); Dietrich et al. (2020), did not focus on the continuous belief updating process (behaviour estimation), or when it focused on behaviour (deceleration) detection (Ackermann et al., 2019) it just relied on a single reaction/detection/identification time point rather than the whole temporal spectrum of an approaching vehicle's scenario. The segmented approach employed in the experiment which was presented in this chapter, appears to have been effective in extracting and measuring these evolving beliefs, as evidenced by the systematic changes and consistent patterns of participants' judgements across segments and a variety of kinematic scenarios (Figure 2.8, Figure 2.9, Figure 2.10). However, a limitation of the experiment was that it was unable to provide any insight on the initial beliefs, i.e., during the time before any of the first segments (first judgement point of each of the 18 scenarios).

It was found that pedestrians were generally able to detect and distinguish the behaviour (maintaining constant speed with no intention to stop or decelerate with the intention to stop) of the approaching vehicle. Different driving manoeuvres and segments (i.e., amount of observation time) were found to influence pedestrians' beliefs (Figure 2.5). This was expected because with more evidence over time and with clearer kinematic distinctions between manoeuvres, pedestrians would become more certain and accurate in judging the approaching vehicle's behaviour. More specifically, the later the pedestrian's belief was

reported, the more certain the reported ratings were towards the correct vehicle behaviour, as also indicated by the statistical analysis. This process could be understood as a form of evidence accumulation, where the input sensory cues (vehicle kinematics) are integrated, over time, to update an internal belief state about the most likely vehicle intention/behaviour (stopping or not stopping). Early on, there was a divergence between the pedestrians' beliefs between deceleration manoeuvre scenarios and the other two manoeuvre scenarios. This divergence is in line with the vehicle kinematics at that stage, i.e., in the first segments of deceleration scenarios, pedestrians were observing kinematic cues indicative of the stopping behaviour, while in the first segments of constant speed and mixed scenarios, those cues were absent. The slight differences between the pedestrians' beliefs in the first segments of the respective constant speed and mixed scenarios could be attributed to the temporal structure of the presented stimuli. The first segments of the mixed scenarios (presenting a constant speed phase) were comparable to later segments of the constant speed scenarios. As a result, the beliefs of these later constant speed manoeuvre segments were similar to the beliefs of the first mixed manoeuvre segments. As segments progressed and the kinematics were separated between mixed and constant speed scenarios, beliefs also diverged. On the other hand, the beliefs observed in later segments of mixed scenarios converged towards the beliefs observed in the later segments of deceleration scenarios. Overall, the pedestrians' beliefs regarding the approaching vehicle's behaviour followed the actual kinematics of the vehicle in a way that one might expect, which provides reassurance that the segment-based experiment design was efficient.

Two interesting findings are related to the mixed manoeuvres, which could be considered as later braking onsets than the respective deceleration manoeuvres (Figure 2.5). First, the average belief's change of direction, from the non-stopping behaviour to the stopping behaviour, was evident and seemed to be happening almost as soon as the vehicle's behaviour changed. Second, after the first time segment the rate with which the average belief of the mixed manoeuvres was increasing towards the stopping behaviour was larger than the rate with which the average belief of the deceleration manoeuvres was increasing towards the stopping behaviour, which aligns with the beliefs-vehicle kinematics relationship that was discussed before. In the later segments of the mixed scenarios the

deceleration rate magnitude was larger than the deceleration rate of the deceleration scenarios and thus the evidence related to the stopping behaviour is stronger in the former case. This observation could validate the finding by Ackermann et al. (2019), that the harder the braking (larger deceleration rate), the faster the identification of the stopping (decelerating) behaviour. However, Ackermann et al. (2019) found no significant influence of the onset of deceleration to the identification of vehicle behaviour, which contrasts with the finding of the current work, where pedestrians were overall more certain that the vehicle was stopping during deceleration manoeuvres than during mixed manoeuvres, despite the lower deceleration rates observed in the former scenarios. This suggests that earlier braking onset, even with a lower deceleration rate, could lead to faster and more accurate beliefs that the vehicle is stopping, possibly due to the longer duration of integrating decelerating cues.

In the cases of larger spatiotemporal distance between the pedestrian and the vehicle, the pedestrians' belief certainty and accuracy were lower (Figure 2.6). The segments, where the discrepancy between the believed vehicle behaviour and the actual vehicle behaviour was largest, were mostly earlier segments with lower vehicle speeds and/or larger TTA, i.e., larger spatiotemporal distances. This observation contrasts with the previous beliefs-vehicle kinematics relationship and suggests that also other factors might influence pedestrians' beliefs, especially when direct perceptual information of deceleration might be less salient. Specifically, the effects of initial speeds and initial TTAs on pedestrians' beliefs remained consistent across all the kinematic scenarios (Figure 2.6). Generally, pedestrians tended to believe more towards the vehicle's stopping behaviour when the initial TTA was longer. Conversely, higher initial speeds led pedestrians to believe that the vehicle was not stopping, which is in line with the findings by Ackermann et al. (2019).

During the experiment there was no auditory feedback (sound cues) related to the approaching vehicle, for example engine and/or tyre noise, braking sounds, etc. That meant that the pedestrian beliefs were based solely on the visual perception of the vehicle's motion. The omission of sound cues was intentional to focus on the interpretation of implicit visual kinematic cues, which aligned with the core mechanism of the suggested model and its reliance on a visual variable. That omission represents a limitation, as sound could influence a pedestrian's

perception of an approaching vehicle and their road crossing behaviour. While some studies have suggested that the addition of vehicle sound did not significantly affect pedestrians' crossing decisions, which were primarily driven by visual cues of speed and distance (Soares et al., 2021), other research has indicated that sound can be important for accurate perception (Wessels et al., 2023). So, it is possible that including realistic sound cues would have influenced the speed and certainty with which the participants updated their beliefs, by providing earlier or more redundant evidence of vehicle deceleration. However, due to the challenges in simulating high-fidelity, context-specific auditory stimuli and the initial focus on visual cues, this influence remains an important area for future investigation.

2.6.2 Computational modelling of pedestrian behaviour estimation

Another contribution of the work presented in this chapter is the development and evaluation of a behaviour estimation model. This model aimed to adapt a Bayesian observer model (Baker et al., 2009), which has previously been successful in capturing humans inferences of others' goals in simplified laboratory tasks, to the road crossing context. To do this, kinematic cues, that researchers have previously proposed are used by humans as interpretation of stopping behaviour, were used. Specifically, the utilised kinematic cues were the deceleration rate (d) of a vehicle and the rate of change of the relative rate of optical expansion ($\dot{\tau}$).

Both behaviour estimation model alternatives suggested (d-based and $\dot{\tau}$-based), were able to provide belief predictions which captured the general trends of average pedestrian beliefs regarding the behaviour of an approaching vehicle. The relatively high Spearman's rank correlations that they achieved, indicate a strong monotonic relationship between their predictions and the average pedestrian beliefs and further suggests their ability to follow the overall patterns of the subjective beliefs. One reason for this relative success could be that humans, like the model, put a lot of emphasis on deceleration cues when inferring whether an approaching vehicle is stopping or not. The proposed model formulates how these cues are processed over time, following an inference process, in which, sensory information input is updating an internal belief state and replicates the beliefs-vehicle kinematics relationship discussed in the previous subsection.

Another key finding of the modelling work presented is the better performance of the $\dot{\tau}$-based alternative in comparison to the d-based one (Figure 2.7). This better belief prediction was expressed both in quantitative (higher Spearman's correlation and lower RMSE between their predictions and the average pedestrian beliefs) and qualitative (fewer poorly predicted datapoints and predicted belief curves closer to the actual average subjective ratings) terms. This provides support for the hypothesis that $\dot{\tau}$ is perceptually available and more useful to pedestrians, rather than observations of the deceleration rate of a stopping vehicle. Another possible reason why the $\dot{\tau}$-based alternative performs better than the d-based alternative is based on the fact that it is situation-adaptive whereas the latter is not. When a vehicle is braking, the value of $\dot{\tau}$ is the same regardless of the specific kinematics of the vehicle, providing a 'universal' value which is indicative of vehicle stopping behaviour. In contrast the deceleration rate d is different for different of stopping manoeuvres. The optimal values for the μ parameters of the likelihood distributions, for the $\dot{\tau}$-based alternative, were found to be -0.5 for the stopping behaviour and -1 for the non-stopping behaviour; in line with the collision avoidance theory (Lee, 1976).

Although, the Bayesian observer model has been very successful in capturing the beliefs of humans in simplified laboratory tasks, the current work suggests that it is somewhat less successful in the traffic setting. Despite achieving high positive correlations, a qualitative comparison revealed the model's inability to capture all observed patterns in human judgments. This failure was most apparent in cases of larger spatiotemporal distance between the pedestrian and the vehicle. An obvious limitation which may be contributing to the model's difficulty to replicate the average pedestrian beliefs, is its reliance on a single deceleration related cue. The empirical observations discussed before, more specifically the influence of initial speed and TTA on beliefs in early segments, suggest that pedestrians, in addition to deceleration related cues, use also other sources of information that affect their behaviour estimation mechanism. Incorporating these sources of information in the modelling of the behaviour estimation mechanisms could possibly provide belief predictions that capture the patterns and nuances of the average pedestrian beliefs better.

A further limitation of the modelling work in this chapter is its exclusive focus on aggregate data. The simplifying assumption of fitting the model to average

participant data to capture general trends in belief updating was adopted. However, this approach does not consider the variability between individuals. A closer inspection of individual participant ratings showed that individuals' belief curves tended to follow the overall patterns of the average subjective ratings. Hence, it could be argued that the model would be limited in fitting the data, exactly, on the per-participant level. Therefore, it seems like the same basic mechanisms indeed may be at play for all participants, such that a model that fits the average data could also probably be fitted to individual participants.

Finally, the values of the behaviour likelihood distribution means (μ_s and μ_{ns}) and their standard deviations (σ_s and σ_{ns}) are not uniquely identifiable because the belief calculation seemed to be depending on the evidence difference of the two possible behaviours. It is likely that if both μ_s and μ_{ns} were shifted equally by a constant value, the resulting belief and correlation would remain the same. An action to correct this non-identifiability would be to fix the values of some of the parameters that seem to have a clear relationship with other parameters.

2.6.3 Summary

This chapter detailed how pedestrian beliefs regarding an approaching vehicle's behaviour (stopping or not) change over time. Different segments (duration of observing the approaching vehicle), driving manoeuvres, vehicle initial speeds and TTAs influenced the pedestrians' beliefs regarding the behaviour of the vehicle. The proposed model, and specifically the $\dot{\tau}$-based alternative, was successful in capturing the general trends of the pedestrians' beliefs, suggesting a basis of modelling of the behaviour estimation mechanism. However, the model did not succeed at capturing all the details and patterns observed in the participants' ratings, especially in scenarios with greater spatiotemporal distances, highlighting areas for model improvement. In those cases, a possible explanation of the nuanced patterns observed in the average pedestrian beliefs might be that these beliefs are based the pedestrians' prior knowledge and expectations.

3 Behaviour estimation through observations and expectations

In this chapter the aim is to propose a more comprehensive behaviour estimation model than the proposed model in Chapter 2, which would also be more successful at predicting the beliefs of pedestrians regarding an approaching vehicle's behaviour. As discussed in Section 2.6, relying solely on observations of a deceleration-related metric is not enough to capture how pedestrians form and update these beliefs. While the vehicle's deceleration rate is important for recognising the driver's/AV's intentions, it does not seem to be the only relevant cue used by pedestrians in estimating the behaviour of the approaching vehicle (Ackermann et al., 2019; Petzoldt et al., 2018). Chapter 2 concluded that when a vehicle is far away (i.e., larger TTA), pedestrians might rely on their prior knowledge and expectations of the current situation they are facing, while on the other hand, it seems that an observation-based behaviour estimation would be more relevant when the vehicle is closer to the pedestrian. As presented in Section 1.6 humans possess self-awareness and also a "theory of mind" regarding other agents with whom they interact. There are distinct psychological theories that elucidate how humans infer the intentions of others by either directly observing their actions (Baker et al., 2009; Pezzulo et al., 2013) or deducing the other agent's behaviour through rational reasoning, or in simpler words expecting the other agent to follow the behaviour that is the most beneficial (value-maximising) to them (Markkula et al., 2023; Jara-Ettinger et al., 2020).

Therefore, in Chapter 3, the more detailed behaviour estimation model that is suggested, is based on the two psychological theories of human intention inference mentioned above. The first will be referred to as observation-based behaviour estimation (Ob) and it is the same type of estimation as in the already proposed model of Chapter 2. The second is based on pedestrians' expectations of what behaviour would be rational or value-maximising for the driver/AV and it will be referred to as value-based behaviour estimation (Vb). Then, a model combining both mechanisms for behaviour estimation will be introduced, referred to here as augmented behaviour estimation model (Ob+Vb). These three models are then evaluated on the same dataset which was collected in the experiment presented in the previous chapter (Section 2.1). To assess the performance of

the Ob, Vb, and Ob+Vb models, and to address the critical issue of model complexity, a Bootstrap Cross-Validation (BSCV) technique (Baker et al. (2009) - where Ob and the experimental design were also inspired from), was employed. Specifically, Ob contained 7 parameters, Vb contained 5, and Ob+Vb contained 12 parameters, making it quite a complex model. This model selection technique was key for mitigating the risk of overfitting, a phenomenon where models fit the training data too closely and fail to generalise to new data. Overfitting is often caused by limited data representation or excessive model complexity, as evidenced by the increasing parameter count. Therefore, BSCV was used to compare the performance and check the generalisability of the models, by artificially expanding the data sample and controlling model complexity.

3.1 Model definition

Building on the computational framework of Markkula et al. (2023), this chapter will progress the modelling of pedestrian beliefs regarding an approaching vehicle's intentions, with the main focus specifically on the mechanisms of behaviour estimation. While Markkula et al. (2023) developed a holistic model of pedestrian crossing behaviour, this work will focus only on and model the pedestrians' belief updating regarding the vehicle's behaviour, allowing for a more detailed investigation of behaviour estimation, which has previously been not addressed in this manner. Specifically, in the current work the Markkula et al. (2023) framework was combined with the Baker et al. (2009) Bayesian observer models of goal estimation. The overall aim remains the same as in Chapter 2, that is to implement models that predict pedestrians' beliefs about the behaviour of an approaching vehicle, mainly expressed as the probability (belief) that the driver/vehicle will decelerate/yield/stop (P_s).

3.1.1 Observation-based evidence

In Section 2.3 a formulation for the observation-based mechanism based on Recursive Bayesian Estimation was introduced. This method involved directly calculating the posterior belief P_b, through a straightforward iterative application of Bayes' theorem, by multiplying the prior belief to the likelihood and then normalising over the sum of the products of the two possible behaviours

(Equation (2.11)). However, this process can be equivalently represented using a normalised exponential (softmax), such that the Bayes update can be seen as just iterative additions to a sum of accumulated evidence.

Specifically, the posterior belief P_b that the vehicle will exhibit behaviour b can be expressed as a normalised exponential, or softmax function of the accumulated observation-based evidence $A_{O,b}$ over the sum of the observation-based evidence for all behaviours b', using the evidence of the previous time step $t-1$:

$$P_b(t) = \frac{e^{A_{O,b}(t-1)}}{\sum_{b'} e^{A_{O,b'}(t-1)}} \tag{3.1}$$

The following observation-based evidence update is also assumed:

$$A_{O,b}(t) = \left(1 - \frac{\Delta t}{T_f}\right) A_{O,b}(t-1) + \frac{\Delta t}{T_u} lnP[\dot{\tau}(t)|\dot{\tau}(t-1), b] \tag{3.2}$$

Where T_f is a time constant which represents how quickly someone forgets older evidence, T_u is a time constant which represents the duration of the evidence updating time step and $P[\dot{\tau}(t)|\dot{\tau}(t-1), b]$ is the probability of the currently observed $\dot{\tau}$ given that the vehicle's current behaviour is b. These probabilities were modelled as normal distributions for the observed $\dot{\tau}$, with mean $\dot{\tau}$ associated to behaviour b and standard deviation σ_b. If it is further assumed that the pedestrian does not forget the observation-based evidence ($T_f \to \infty$) and the duration of the evidence updating time step is set to be equal to the model time step ($T_u = \Delta t$), then the Recursive Bayes Estimation formulation (Equation (2.11)) can be derived:

$$P_b(t) = \frac{e^{A_{O,b}(t-1)}}{\sum_{b'} e^{A_{O,b'}(t-1)}} \overset{(3.2)}{\Longrightarrow}$$

$$P_b(t) = \frac{e^{A_{O,b}(t-1)} P[\dot{\tau}(t)|\dot{\tau}(t-1), b]}{\sum_{b'} e^{A_{O,b'}(t-1)} P[\dot{\tau}(t)|\dot{\tau}(t-1), b']}$$

$$= \frac{\frac{e^{A_{O,b}(t-1)}}{\sum_{b''} e^{A_{O,b''}(t-1)}} P[\dot{\tau}(t)|\dot{\tau}(t-1), b]}{\sum_{b'} \frac{e^{A_{O,b'}(t-1)}}{\sum_{b''} e^{A_{O,b''}(t-1)}} P[\dot{\tau}(t)|\dot{\tau}(t-1), b']}$$

$$= \frac{P_b(t-1) P[\dot{\tau}(t)|\dot{\tau}(t-1), b]}{\sum_{b'} P_{b'}(t-1) P[\dot{\tau}(t)|\dot{\tau}(t-1), b']}$$

The derivation above provides the special case of pure Bayesian observation-based behaviour estimation. Initially, the pedestrian does not have any prior $\dot{t}$ observations available. Thus, all $A_{O,b}$ are equal to zero and the beliefs regarding the two possible vehicle behaviours (decelerating/stopping and maintaining constant speed/not stopping) are equal:

$$P_s(t=0) = \frac{e^{A_{O,s}(t-1)}}{e^{A_{O,s}(t-1)} + e^{A_{O,ns}(t-1)}} = \frac{e^0}{e^0 + e^0} = 0.5$$

$$P_{ns}(t=0) = \frac{e^{A_{O,ns}(t-1)}}{e^{A_{O,ns}(t-1)} + e^{A_{O,s}(t-1)}} = \frac{e^0}{e^0 + e^0} = 0.5$$

3.1.2 Value-based evidence

As mentioned in the introduction of Section 3.1, a pedestrian might form their beliefs regarding the approaching vehicle's behaviour based on what behaviour they would expect to be the most rational or value-maximising for the AV or the driver of the vehicle. Affordance theory posits that perception is not only the interpretation of someone's surroundings but is also used to identify potential interactions and enable the pursuit of effective courses of action (Chemero, 2003; Lio et al., 2020). Building upon the affordance theory statement above and the principles of reward-driven behaviour, a reward function was adopted, which models how a driver could maximise the value of their behaviour over a prediction horizon, as has been presented in Markkula et al. (2023) and based on models of optimal human motor and locomotor control (Gawthrop et al., 2011; Hoogendoorn and Bovy, 2003; Wang et al., 2015).

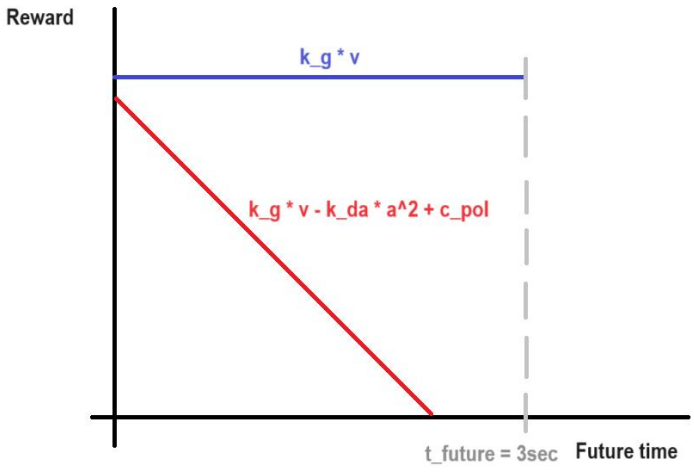

Figure 3.1 - Example of future rewards of the two possible behaviours that the vehicle can exhibit

It was assumed that pedestrians believe that the driver's behaviour and actions are guided by a core reward function, designed to balance the desire for efficient progress towards their goal with the discomfort of braking and the collaborative consideration of yielding. This function is defined as:

$$R(t) = k_g v(t) - k_{da} a^2(t) + c_{pol} \quad (3.3)$$

Here, $R(t)$ represents the reward at time t, quantifying the desirability of a certain behaviour. k_g is a reward parameter associated with speed, $v(t)$ is the vehicle's current speed, k_{da} is a cost parameter associated with deceleration, $\alpha(t)$ is the vehicle's current acceleration, and c_{pol} is a constant representing a politeness reward. Equation (3.3) effectively captures the driver's trade-offs:

- The positive term $k_g v(t)$ encourages driving at higher speeds, reflecting the desire to reach the destination as fast as possible.
- The negative term $-k_{da} a^2(t)$ penalises acceleration/deceleration, representing the discomfort of accelerating/braking, especially in cases where the vehicle changes its speed abruptly.
- The positive term c_{pol} introduces a reward for yielding in the presence of pedestrians at the curb, reflecting the driver's anticipation of how their behaviour might be perceived by a naive observer, which can be either cooperative (higher reward) or non-cooperative (no reward). The use of "politeness" and "pro-social behaviour" definitions in this thesis require further clarification, as the approaching vehicle in the experiment was an AV without a visible human driver. The application of these terms was intentional, serving to model the expectations of the human pedestrian based on the following assumption. Pedestrians possess a Theory of Mind of how road users interact, which is based on their lifelong experience with manually driven vehicles. In interactions with manually driven vehicles the act of yielding is considered pro-social behaviour which indicates cooperation, prioritising pedestrians' safety and convenience. When interacting with an AV, pedestrians try to infer the vehicle's intent by projecting the politeness social value onto its behaviour (Lanzer et al., 2020; Ribino, 2023; Tsui et al., 2010). The politeness term formalises the pedestrians' expectation that a rationally operating vehicle will derive a positive social reward from acting cooperatively (i.e., yielding).

This unified reward function can be applied to analyse the two possible behaviour horizons: maintaining constant speed with no intention to stop and decelerating with the intention to stop.

In the non-stopping behaviour horizon, which is illustrated as the blue curve in Figure 3.1, the driver maintains a steady velocity, resulting in zero acceleration ($\alpha(t) = 0$). Additionally, in this case, it is assumed that the politeness reward is inactive, therefore $\mathrm{c_{pol}} = 0$. Substituting these values into Equation (3.3), the simplified reward function for maintaining constant speed/not stopping is obtained:

$$R_{ns}(t) = k_g v(t) \tag{3.4}$$

Conversely, in the stopping behaviour horizon, which is illustrated as the red curve in Figure 3.1, the driver reduces speed to stop before the pedestrian's position, leading to a non-zero deceleration rate, specifically $\alpha(t) = {a_{req}}^2(t)$, where a_{req} represents the required deceleration rate to come to a complete stop at a safe distance from the pedestrian. In this case, the politeness term c_{pol} remains active. Substituting these into Equation (3.3), the reward function for stopping is obtained:

$$R_s(t) = k_g v(t) - k_{da} {a_{req}}^2(t) + c_{pol} \tag{3.5}$$

Notably, the politeness constant reward could instead be formulated as a selfishness constant cost in Equation (3.4). The required deceleration rate was calculated using the deceleration formulae:

$$a_{req}(t) = \frac{v(t)}{t_{stop}} \tag{3.6}$$

$$S = \frac{1}{2} a_{req}(t)\, t_{stop}^2 \tag{3.7}$$

Here, $v(t)$ is the current vehicle speed, t_{stop} is the time needed for the vehicle to come to a stop and S is the distance that the vehicle will travel during the deceleration manoeuvre and is equal to $D(t) - D_{stop}$, $D(t)$ being the current distance between the vehicle and the pedestrian and D_{stop} being the distance between the vehicle and the pedestrian when the vehicle has fully stopped. In accordance with the "The Highway Code" (2023), D_{stop} was chosen to be approximately 2 to 2.5 m.

To calculate the value-based evidence $A_{V,b}$ of behaviour b:

$$A_{V,b}(t) = \int_{t}^{t+T_h} R_b[x(t)]dt' \tag{3.8}$$

Here t is the current time, $T_h = 3\ s$ is the time horizon that was chosen, as it was equal to the smallest initial TTA and scenario duration of the experiment described in Chapter 2 and $R_b[x]$ is the reward value for behaviour b and current kinematic state of the vehicle x. Essentially, $A_{V,b}(t)$ can be found by calculating the area underneath the R_b curves in Figure 3.1.

3.1.3 Pedestrian beliefs of the vehicle's behaviour

The distinction between the inputs for the Ob and Vb models is based in the level of cognitive processing assumed, and not on whether the kinematic variable is physically observable. All kinematic cues (speed, distance, acceleration) are technically observed variables. However, the Ob component was focused exclusively on the perceptually salient cue for predicting the behaviour of the approaching vehicle, in this case the $\dot{\tau}$. This represented a perceptual processing mechanism that updates beliefs based on immediate sensory evidence. Conversely, the Vb component used observed variables like speed and distance as inputs to a utility calculation (i.e., expected reward). This process models the pedestrian's rational expectation that the perceived driver/AV acts to maximize its reward by balancing objectives like minimising braking cost and maximising maintaining speed towards a goal. Thus, Ob used direct evidence of motion change, while Vb used observable kinematic cues to form reward expectations of rational behaviour. Finally, the pedestrians' overall goal (e.g., to cross safely) is necessary information for the road crossing decision but is not required as an input for the intermediate task of behaviour estimation modelled here.

If the observation and value-based behaviour estimation mechanisms are combined, then the overall evidence that the vehicle is exhibiting behaviour b, could be modelled as the weighted sum of the estimated observation-based evidence $A_{O,b}$ and value-based evidence $A_{V,b}$.

$$A_b(t) = \beta_O A_{O,b}(t) + \beta_V A_{V,b}(t) \tag{3.9}$$

$$P_b(t) = \frac{e^{A_b(t-1)}}{\sum_{b'} e^{A_{b'}(t-1)}} \tag{3.10}$$

Where β_O and β_V are the weights of the respective evidence type.

As stated in Subsection 3.1.1, the model that was suggested in Chapter 2, was reformulated to operate on evidence rather than probabilities and its belief predictions were successfully replicated by the evidence accumulation formulation. However, it is important to first consider the theoretical motivation for the use of the softmax application in the model. As detailed in Markkula et al. (2023), the softmax function was used for 3 important reasons: (1) it allowed for the representation of Bayesian belief update as additive evidence accumulation, (2) it has been commonly used to model decision probabilities as a function of rewards, particularly in behavioural game theory (Wright and Leyton-Brown, 2017) and (3) it provided a straightforward framework for integrating observation-based and value-based mechanisms within a single model. In addition, the correlational analysis in Chapter 2 (Figure 2.7) showed that the relationship between the model predictions and the average subjective ratings is not linear, but rather it seemed that an S-shaped correlation could describe the relationship better. Many natural processes, like the learning curves of complex systems, start slowly, accelerate, and then level off; this seems to be resembling the overall data trend in Figure 2.7. In cases where a precise mathematical representation is unavailable, an S-shaped function has often been applied to approximate this pattern (Gibbs and MacKay, 2000). More specifically, borrowing from deep learning (Goodfellow et al., 2016) and pattern recognition (Bishop, 2006), the softmax function has also been frequently used to map a vector of feature variables to a posterior probability distribution.

Despite these theoretical advantages, early tests of the model presented in this chapter suggested that the standard softmax function was not capable of fully capturing the precise shape of the relationship between model predictions and the average subjective ratings reported by human observers in the experiment. To better align the model predictions with the observed subjective ratings, a modified version of the softmax transformation function was required. Drawing inspiration from the Richards' family of growth-models (Richards, 1959; Tjørve and Tjørve, 2010), modifications were introduced to Equation (3.10). It is important to clarify the relationship between the behaviour probabilities generated by the model, $P_b(t)$, and the judgments provided by the participants during the experiment, as described in Subsection 2.1.4. While the output of the models is probabilities representing the pedestrians' beliefs regarding the behaviour of an

approaching vehicle, participants provided subjective ratings. These ratings were subsequently transformed into belief scores to facilitate comparison with the model's probabilistic output. The introduction of parameters B and M (see Equations (3.11) and (3.12)) addresses the potential discrepancy between the model's probabilistic output and the subjective ratings provided by participants. Although the softmax function transforms evidence to probabilities, the relationship between belief probabilities and human confidence ratings is not necessarily linear. This modification achieved a linear relationship between the model predicted beliefs and the average subjective ratings, addressing the observed non-linearity between the two, as will be seen in Section 3.4.

$$P_s(t) = \frac{e^{B(A_s(t-1)+M)}}{\sum e^{B(A_{ns}(t-1))} + e^{B(A_s(t-1)+M)}} \tag{3.11}$$

$$P_{ns}(t) = \frac{e^{B(A_{ns}(t-1))}}{\sum e^{B(A_{ns}(t-1))} + e^{B(A_s(t-1)+M)}} \tag{3.12}$$

Where P_s and P_{ns} are the predicted beliefs for the decelerating and constant speed behaviour respectively, A_s and A_{ns} are the evidence of the decelerating and constant speed behaviour, respectively. B is an evidence scaling factor which adjusts the sensitivity of the softmax function to changes in accumulated evidence, thereby enabling the model to better capture the observed variability in subjective belief ratings across different scenarios and M is an added offset constant applied exclusively to the decelerating behaviour evidence, A_s, accounting for potential biases or baseline differences in participants' perception of the decelerating behaviour. Overall, this transformation facilitates a more accurate representation of how participants translate internal beliefs into subjective ratings.

The parameter B, which is referred to as the inverse temperature constant in literature, is a growth-rate constant that controls the slope at the inflection point and thus the overall shape of the S-shaped transformation function. Referring back to Equation (3.4) and reiterating the assumptions of $T_f \to \infty$ and $T_u = \Delta t$, it becomes evident that B has the reverse effect of T_u. Consequently, lower values of B, correspond to longer duration of evidence-updating. Through the lens of reinforcement learning, for values of $B \to 0$, both behaviours would have nearly the same probability. Conversely, as B increases, the influence of accumulated evidence values on the resultant probability also increases. The impact of the

scaling parameter B on the predicted beliefs is visualised in Figure 3.2. This figure illustrates the mapping from behaviour evidence to the belief that the approaching vehicle is stopping (P_s) or not (P_{ns}) for the participants with examples of a low and a high B value. As shown, a lower B value (blue curve) results in a more gradual shift in belief as the evidence changes, indicating a lower sensitivity to the accumulating evidence. Conversely, a higher B value (orange curve) leads to a much steeper transition, suggesting that individuals with a high B are more decisive and their beliefs change more rapidly with even small changes in the perceived evidence.

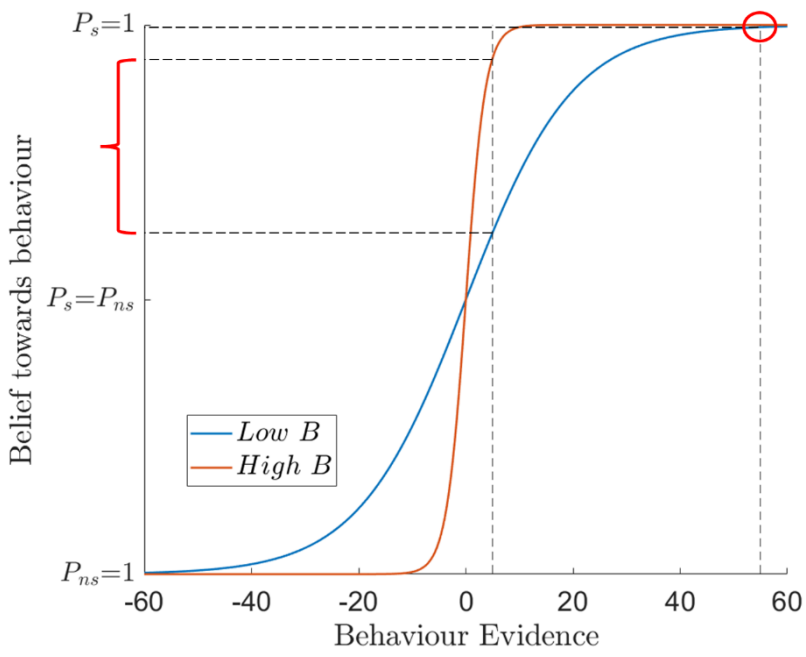

Figure 3.2 - Effect of scaling parameter B on pedestrians' beliefs

The parameter M, which is referred to as soft margin in the literature, is an added constant that has been used in classification tasks to reduce intra-class and increase inter-class separation, by introducing a distance margin into the logits. In the current application, this distance margin could be translated as added evidence to (or bias towards) only one of the two possible behaviours: in this case the decelerating behaviour. Thus, M is affecting the vertical position of the predicted beliefs and acts like an additional prior belief in the modified softmax transformation function. So, the lower the value of M, the lesser the bias towards the decelerating behaviour. The model lacks any other parameter that explicitly represents this form of prior belief. Consequently, the incorporation of M does not only serve to align the beliefs predicted by the model with empirical data. Rather,

it is underpinned by the theoretical concept of equiprobable prior beliefs – how people perceive and assign probabilities to events, especially when faced with uncertainty. The cognitive bias called equiprobability bias is the tendency for people to assume that random events are equally likely, even when there is no logical basis for that assumption (Gauvrit and Morsanyi, 2014; Lecoutre, 1992). Essentially, people often think that everything has an equal chance. In ideal probabilistic reasoning, especially in Bayesian statistics, prior beliefs are initial assumptions about the likelihood of events. Some theoretical models use equiprobable priors as a neutral starting point (especially when there is genuine lack of information); as was also assumed in the purely observation-based behaviour estimation. However, in reality, people's prior beliefs are heavily influenced by experience, knowledge, personal biases and emotional factors (Kapons and Kelly, 2023). Whether or not people lean towards equiprobable beliefs depends heavily on the context. For example, in games of chance with known, fair systems (like a fair die), people might correctly assume equiprobability, but in real-world situations with complex, uncertain factors, people's beliefs are rarely equiprobable (Tversky and Kahneman, 1974). So, it is reasonable to assume that pedestrians do not possess equiprobable prior beliefs, and M could represent this inherent bias.

Notably, even though B and M affect the absolute values of the belief probabilities, they have no effect on the rank order of the predicted beliefs. This is because linear transformations like scaling and shifting do not affect rank correlations, which are based on the ordinal relationship between data points.

3.1.4 Model parameters

Table 3.1 lists the free parameters associated with purely Ob, purely Vb, and augmented Ob+Vb models. The search ranges presented in the table reflect the specific intervals of values within which the fitting algorithm, which will be described in detail in Section 3.2, is allowed to search for the optimal parameter values that best match the data.

Table 3.1 - Free parameters' search ranges

Parameters	Range	Ob	Vb	Ob+Vb
k_g	[0, 1]		+	+
k_{da}	[0, 1]		+	+
c_{pol}	[0, 1]		+	+
μ_c	[-1, 0]	+		+
μ_d	[-1, 0]	+		+
σ_c	[0, 1]	+		+
σ_d	[0, 1]	+		+
β_O	[0, 1]			+
β_V	[0, 1]			+
T_f	[0, 100]	+		+
B	[0, 1]	+	+	+
M	[0, 100]	+	+	+

3.2 Model fitting

In order to fit the parameters of the three behaviour estimation models, the same data that were collected during the experiment, described in Section 2.1, were used. The free parameters of the three models, as well as their respective research ranges, are presented in the subsection above. In Chapter 2 a grid search fitting method was described, which was used in accordance with the Baker et al. (2009) approach. However, the total number of the parameters has increased, and fitting the augmented model could be computationally incredibly costly. Thus, there was a need to use a model fitting method which would be able to decrease the computational cost and running times significantly for a higher number of parameters. For that reason, CRADLE – a combined local and global derivative-free optimisation algorithm (Leonetti et al., 2012) was used[3]. This technique has found application in policy optimisation (Leonetti et al., 2013) and system identification (Karras et al., 2013). This optimisation algorithm, which is a modification of Price's optimisation method (Price, 1983), combines a stochastic

[3] The specific implementation of the optimisation algorithm (CRADLE), which was used in the current thesis for model fitting, was a MATLAB adaptation developed by the first author of the original paper (Leonetti et al., 2012), Dr Matteo Leonetti.

global search with a deterministic local search. Its global nature makes it particularly well-suited for dynamic decision-making problems which need to be solved on-line and where a priori knowledge of a good initial solution is unavailable.

The following optimisation problem is considered: minimising a cost function $J(\theta)$ over a convex domain D in n-dimensional space. In the current case, following the Baker et al. (2009) approach also adopted in Chapter 2, this cost function is the Spearman's rank correlation between the model's belief predictions given a set of parameters θ and the participants' average subjective ratings. However, since higher correlations indicate better model fit, the negative correlation as the cost function to be minimised was used. For this optimisation problem the set of parameters θ would include all parameters except B and M, as they can be omitted due to the fact that they do not affect the overall rank correlation, as stated in Subsection 3.1.3.

As a reminder from Chapter 2, the Spearman's rank correlation coefficient, denoted as ρ_S, is calculated based on the ranks of the data, not the raw values. Specifically, for two sets of data, X (model predicted beliefs) and Y (participants' average beliefs), and their respective ranks are RX and RY ($d_i = RX_i - RY_i$):

$$\rho_S = 1 - \frac{6\sum d_i^2}{n(n^2-1)} \tag{3.13}$$

From Equations (3.11) and (3.12), it can be seen that the parameter M adds a constant to the evidence A_s only in the numerator of the softmax function. Since the rank correlation is based on the order of the predicted beliefs, adding a constant to all values associated with the decelerating behaviour (P_s) will shift these values up, but it will not alter their relative ranking. Therefore, M does not affect the d_i values in the rank correlation calculation, and consequently, it does not affect ρ_S.

Similarly, the parameter B multiplies the evidence A_b in both the numerator and the denominator of the softmax function. While multiplying the evidence by a constant might affect the absolute values of the predicted beliefs, it will not change their relative ordering. Therefore, from a purely mathematical perspective, B could also be omitted from the parameter set θ, as its multiplication to the evidence is happening on both the numerator and the denominator of the

softmax fraction (Equations (3.11) and (3.12)), and similarly to M, the overall rank correlation is not affected. However, due to computational limitations, values of $B \cdot A_b$ extremely close to each other, i.e., having a difference lower than MATLAB's numerical precision, led to those values being recognised as equal, eventually affecting the overall ranking. For that reason, the parameter B was included in the set of parameters for the optimisation problems.

As previously discussed, the model provides probability outputs (beliefs), while participants provided subjective ratings. While a Spearman's rank correlation metric assesses the model's ability to predict the order of participants' beliefs, an RMSE would evaluate the model's accuracy in predicting the magnitude of those beliefs. So, even though the Baker et al. (2009) fitting method used the correlation between the model's predicted beliefs and the average subjective ratings of the participants, it was decided to consider another optimisation problem where the minimised cost function this time would be the root mean square error (RMSE) between the model predicted beliefs and the average subjective ratings. The RMSE was used to measure the average difference between them, with lower RMSE values indicating more accurate model predictions (an RMSE = 0 would mean that the model predicted beliefs match perfectly the average pedestrian beliefs).

$$RMSE = \sqrt{\frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{n}} \tag{3.14}$$

Here, n is the total number of observations (in this case the 72 segments), y_i is the average pedestrian belief value for the i^{th} observation and $\hat{y}_i$ is the model predicted belief value for the i^{th} observation.

Eventually, two fitting methods were tested:

- The first method included two steps. The objective of the first step was to maximise the correlation between the model predictions and the average subjective ratings and obtain the parameter set that produced that correlation. For the second step, using the parameters obtained in the first step, the objective was to minimise the RMSE between the model's predicted beliefs and the average subjective ratings, which would provide the best-fitted M parameter value.

- The second method included only one step, and the objective was to minimise the RMSE between the model predictions and the average subjective ratings and obtain the whole parameter set, that produced that RMSE.

After testing both fitting methods two main takeaways were drawn: a) the second method was more computationally intensive than the first one, b) the second method provided parameter sets that not only resulted in the lowest RMSE values but also produced the highest rank correlations. Conversely, the first method provided parameter sets that resulted in the highest rank correlations but not necessarily the lowest RMSE values. So, even though the first fitting method was faster, the second fitting method seemed more appropriate since it provided a better alignment in terms of both order and magnitude between the model predictions and the average pedestrian beliefs.

3.3 Model selection

To compare the predictive accuracy of the models fitted using the methods described earlier, a quantitative approach was employed. To assess the differences in predictive accuracy between the three models, bootstrap cross-validation (BSCV; Baker et al., 2009) was utilised. BSCV is a non-parametric robust technique for model selection that assesses the goodness-of-fit, while mitigating overfitting and accounting for model complexity by evaluating model fit on multiple resampled datasets.

As illustrated schematically in Figure 3.3, the BSCV analysis involves a large number of iterations, where random subsets of participants' data are selected to train models (sampled uniformly with replacement). For each iteration, the model parameters are optimised to maximise goodness of fit (in this case to maximise correlation or minimise the RMSE) using the sampled training data, and the resulting fitted model is then evaluated on a complementary testing dataset, consisting of the data that was not sampled into the training subset. The ranges of goodness-of-fit values (correlation or RMSE) achieved by each model across the 1000 iterations are presented as distributions of model goodness-of-fit which were obtained from these repeated model optimisations. These distributions were

tested to compare predictive performance of the three models. Better predictive performance is indicated by distributions leaning towards higher correlation or lower RMSE values. While minimally or non-overlapping distributions indicate more substantial differences, overlapping distributions suggest that differences in model performance may not be statistically significant. The width of the distributions represents the variability of the performance of each model across the different BSCV data subsets, with wider distributions indicating greater sensitivity to training data. In this case, each BSCV analysis (one for each model) utilised 1000 iterations, sampling 10 participants (from the total of 30) for each training subset.

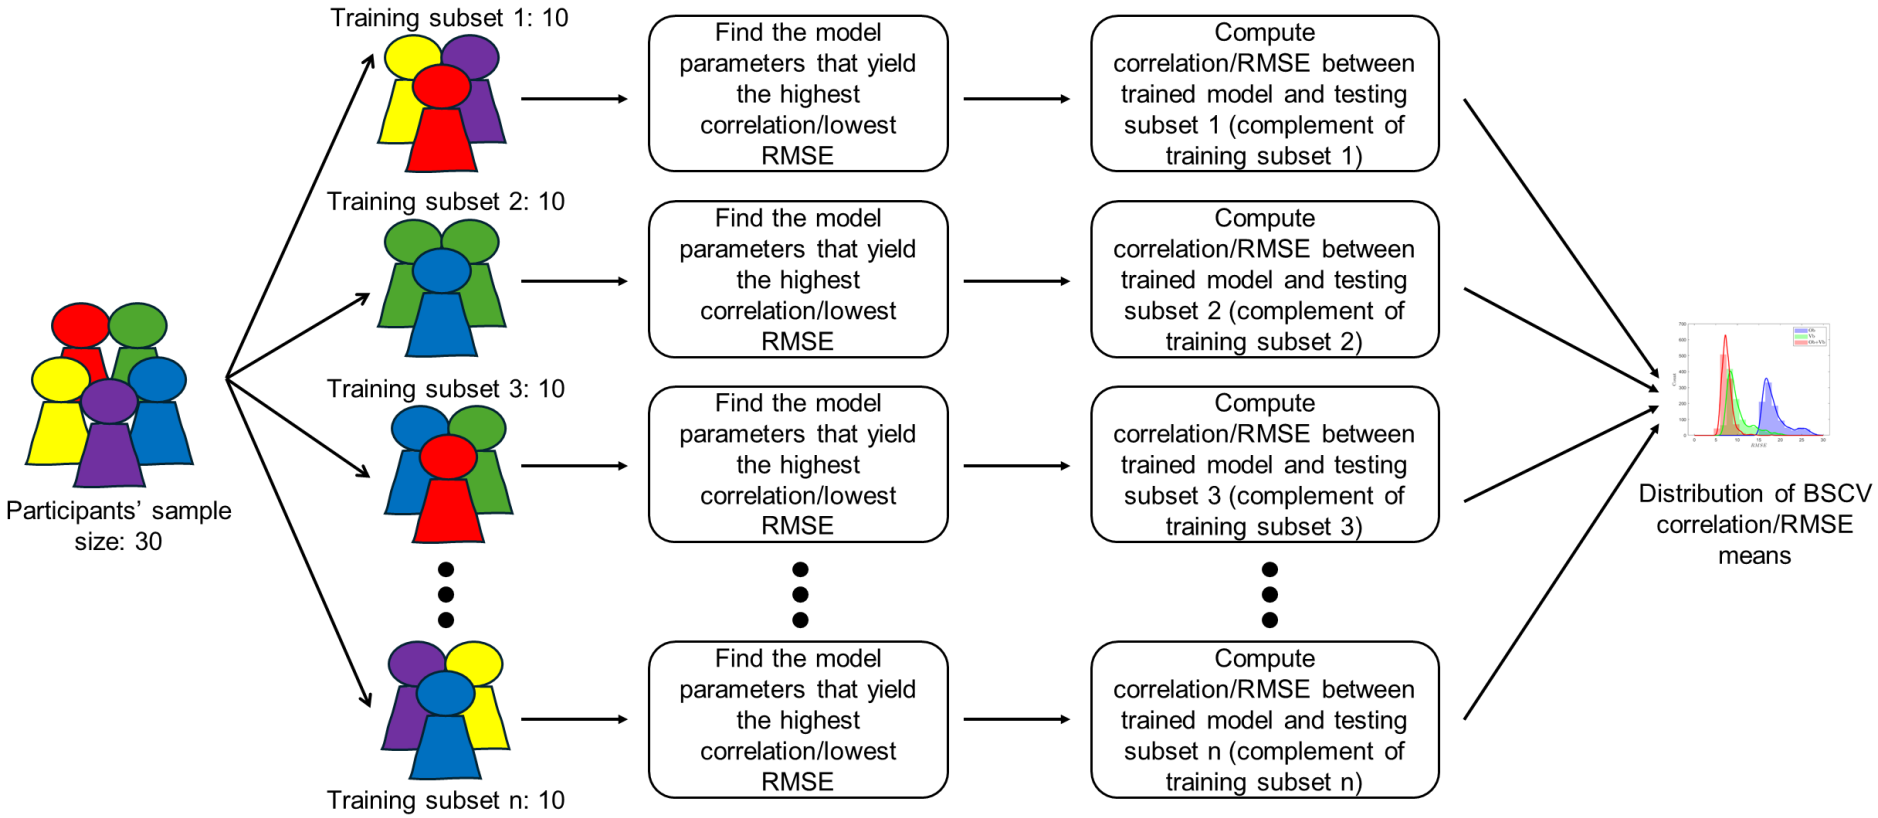

Figure 3.3 - Bootstrap cross-validation (BSCV) technique for model selection

3.4 Results

In this section results from several analyses are presented, to show how accurately the suggested models (Observation based – Ob, Value based – Vb and Observation and Value based – Ob+Vb) predicted pedestrians' beliefs regarding the approaching vehicle's behaviour, which were collected as subjective judgements in the experiment described in Section 2.1.

3.4.1 Model predictions vs subjective ratings

Firstly, the analyses focused on how accurately the three different models predicted pedestrians' average subjective beliefs by using both measures of correlation (Spearman's rank correlation) and error (RMSE). The model predictions presented in this subsection have been obtained by using the second model fitting process (minimisation of RMSE between the model predictions and the average subjective beliefs of the pedestrians), out of the two that were described in Section 3.2, for the full sets of parameters. The best-fitting parameter settings are presented in Table 3.2.

Table 3.2 - Best-fitting parameter settings

	Ob	**Vb**	**Ob+Vb**
k_g		0.88	0.57
k_{da}		0.36	0.51
c_{pol}		0.83	1
μ_{ns}	-1		-1
μ_s	-0.47		-0.75
σ_{ns}	0.55		0.56
σ_s	0.57		0.58
β_O			0.32
β_V			0.53
T_f	4.56		100
B	0.17	0.14	0.20
M	0	16.13	6.26

To visually represent the performance of these models with the best-fitting parameters listed in Table 3.2, Figure 3.4 displays scatter plots of the model predicted beliefs compared to participants' average ratings. As shown in Figure 3.4, Vb and Ob+Vb exhibit the fewest poorly predicted datapoints, while Ob has the most. Similarly, in terms of correlation, Vb and Ob+Vb were able to produce values of Spearman's rank correlation almost equal to 1, showing an almost

perfect positive monotonic rank association between pedestrians' average beliefs and model predictions, while Ob showed a weaker correlation. The fact that Vb's and Ob+Vb's scatter points are so closely gathered to the identity line (Pearson's rho equal to 1, in terms of correlation), means that Vb's and Ob+Vb's predictions are almost identical to the pedestrians' average beliefs (almost perfect positive linear relationship). Lastly, RMSE is the error metric that shows how close the scatter points are to the identity line, so the lower the RMSE, the better the model predictions. From Figure 3.4, Ob+Vb had a lower RMSE than Vb, while Ob was the worst performing model in terms of error.

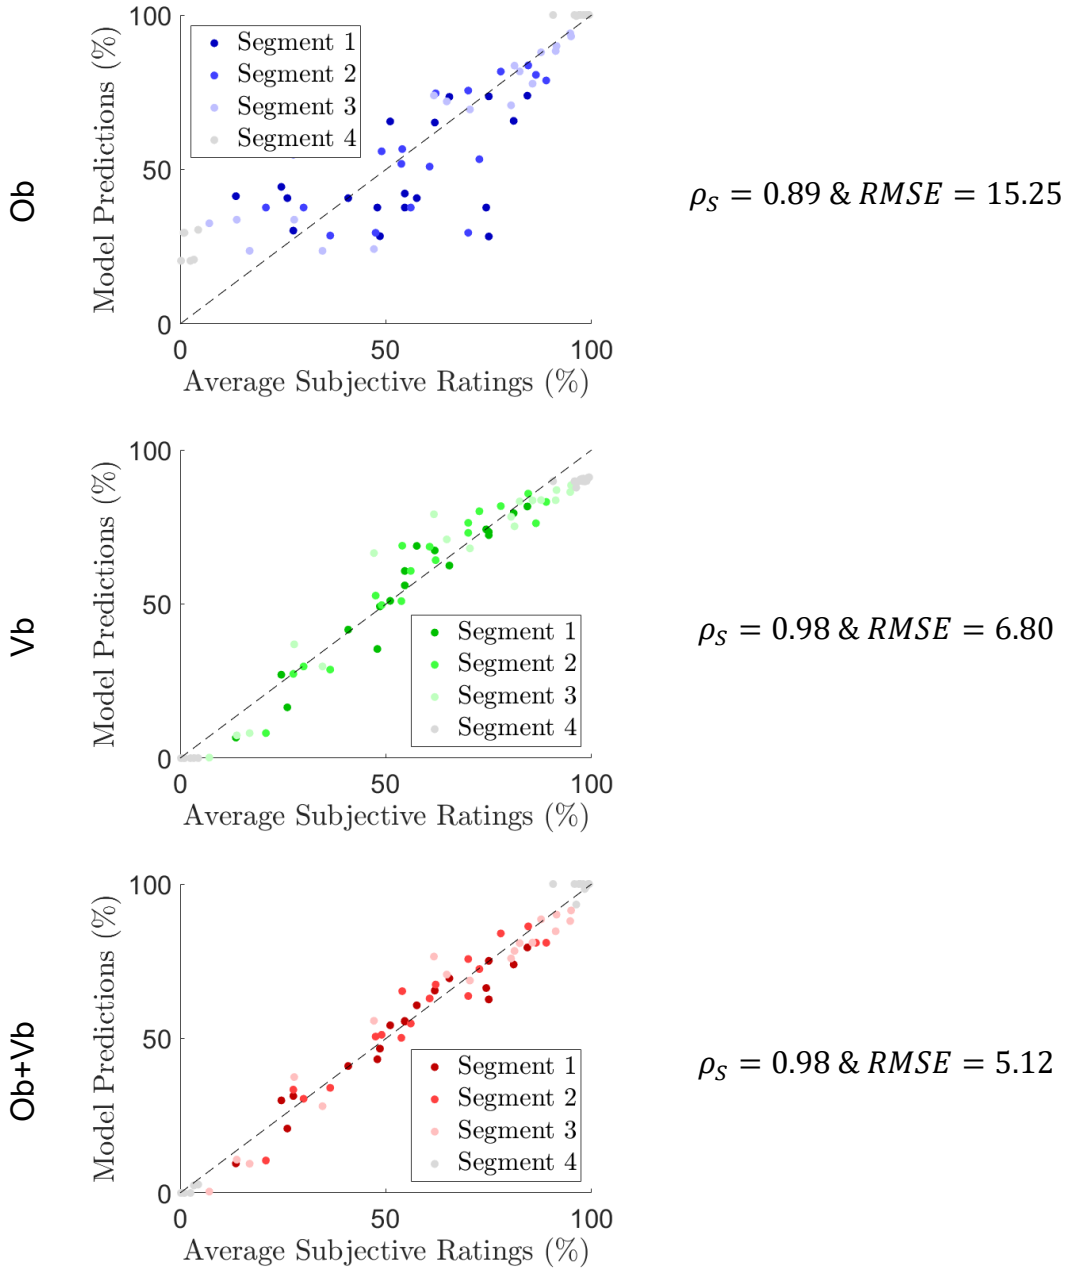

Figure 3.4 - Scatter plots of model predictions using best-fitting parameter settings (y-axes) versus pedestrians' average beliefs (x-axes) for all 18 kinematic scenarios of the approaching vehicle

The following analysis compares the average pedestrian beliefs and the model predictions for all 18 kinematic scenarios (as a reminder the details of these kinematic scenarios are illustrated in Figure 2.2). This analysis shows how different vehicle kinematic conditions affect the pedestrians' belief-updating process regarding the approaching vehicle's behaviour and shows how closely the suggested models captured the average of pedestrians' beliefs.

The different vehicle manoeuvres are highlighted in the following three figures: Figure 3.5 (constant speed), Figure 3.6 (deceleration) and Figure 3.7 (mixed). Each is divided into the three different initial vehicle speeds (rows) and the two different initial TTAs (columns), resulting in 6 different kinematic scenarios per vehicle manoeuvre. All 18 plots illustrate several general patterns of pedestrians' beliefs predicted by the three different models. In Figure 3.5, Figure 3.6 and Figure 3.7 the predictions of Ob, Vb and Ob+Vb are presented by the blue, green and red curves, respectively, while the average pedestrians' beliefs are illustrated as black standard error of the mean bars at the predetermined judgment timings (as described in Section 2.1).

In the constant speed manoeuvres (Figure 3.5), Vb and Ob+Vb are performing better than Ob in capturing the patterns of the pedestrians' beliefs. At the beginning of each scenario, Ob always predicts that pedestrians are completely uncertain between the two possible vehicle behaviours (50% belief). As can be seen in Table 3.2, Ob had its M parameter fitted to zero, meaning that the optimisation algorithm found that no prior shift was optimal for this model, which contributed to this consistent 50% initial belief. Conversely, Vb and Ob+Vb were able to predict that early beliefs would not be 50-50% but would likely be affected by the speed and TTA of the vehicle. Another limitation of the Ob model is that it is incapable of providing different beliefs between the six different constant speed scenarios, since there is no difference in the observed tau-dot (always equal to -1). That limitation is addressed in Vb and Ob+Vb, which make use of other sources of information beyond tau-dot observations. The greater the vehicle's current speed, the more Vb's and Ob+Vb's predicted beliefs lean towards the non-stopping behaviour (i.e., towards $P_s = 0\%$), in line with the beliefs of human participants. Then, the larger the vehicle's TTA, the lower the required deceleration, and eventually, the less Vb's and Ob+Vb's predicted beliefs lean towards the non-stopping behaviour. The patterns described above can be further

understood by examining the different sources of evidence (more details in Subsection 3.4.2).

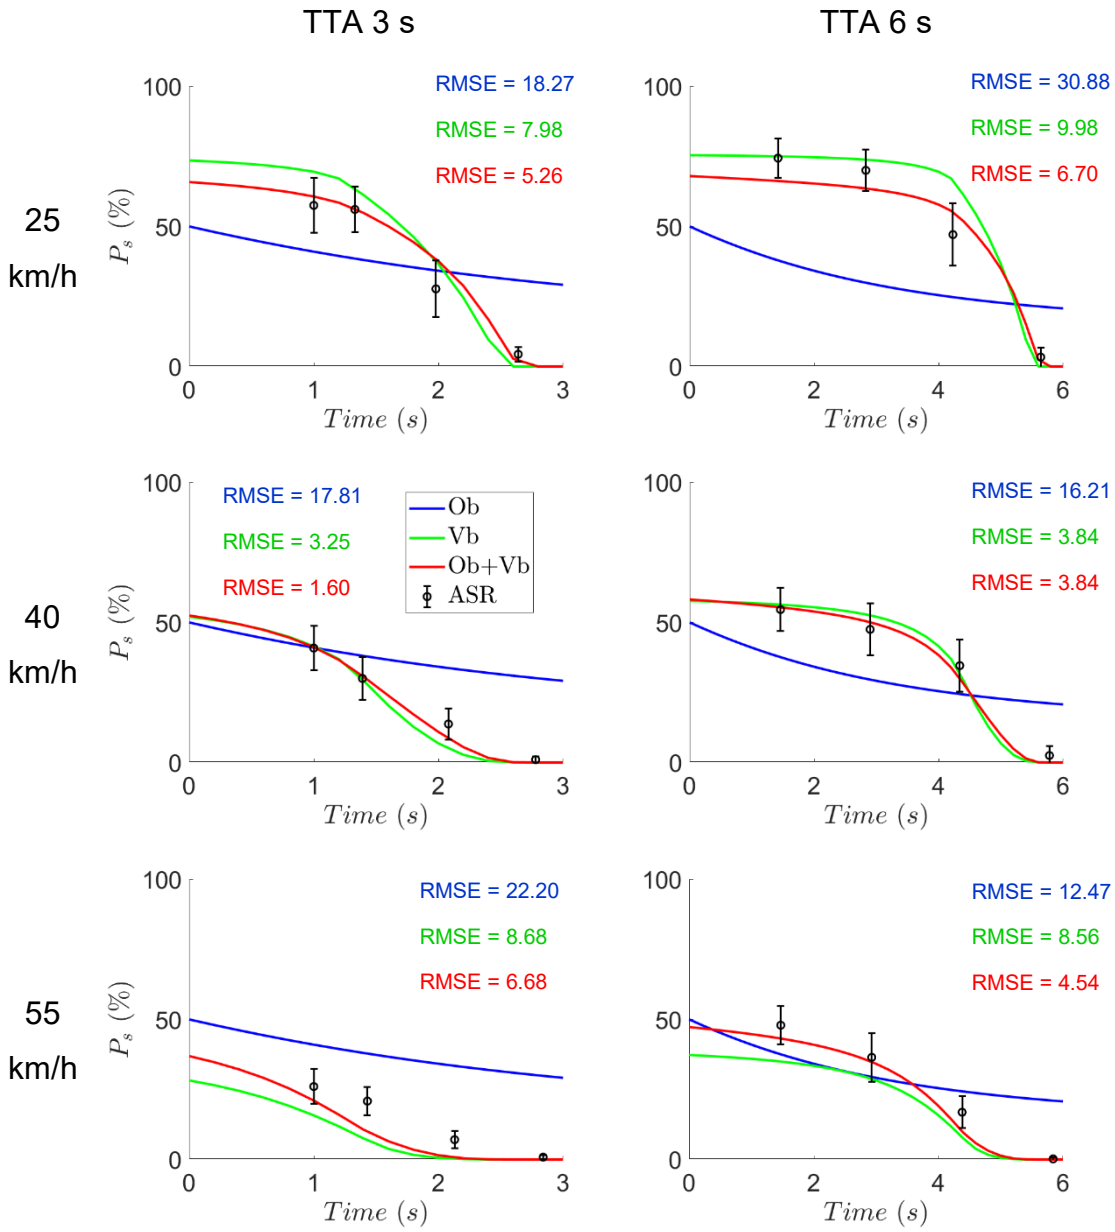

Figure 3.5 - Comparison between model predictions and pedestrians' beliefs for the constant speed manoeuvres

In the deceleration manoeuvres (Figure 3.6), Ob+Vb is performing better than Ob and Vb in capturing the patterns of the pedestrians' beliefs. Ob's initial predictions are unable to capture the effects of the different vehicle speeds and TTAs on the pedestrians' beliefs. However, as time passes and $\dot{\tau}$ values change, Ob's later predictions are more accurate. Vb's initial predictions are capturing the speed and TTA effects on the pedestrians' beliefs, but its final segments' predictions are not that accurate. Thε Vb's limitation happens because the evidence of the two possible behaviours are equal in the end of each deceleration

manoeuvre (more details in Subsection 3.4.2). Ob+Vb's predictions are benefitted by Vb early and by Ob later.

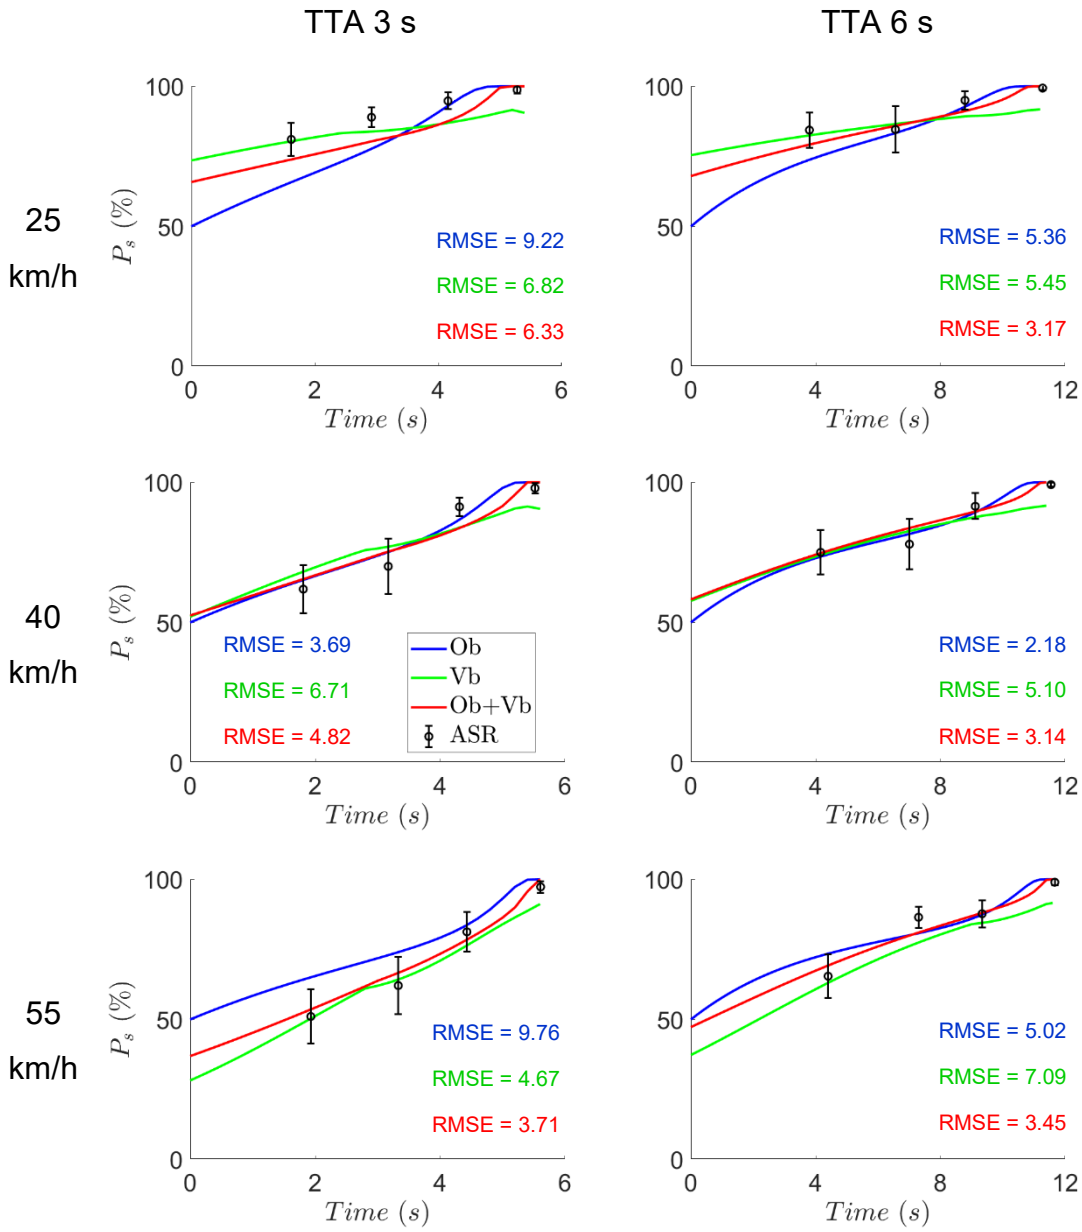

Figure 3.6 - Comparison between model predictions and pedestrians' beliefs for the deceleration manoeuvres

In the mixed manoeuvre scenarios (Figure 3.7), it can be seen that Ob+Vb is again performing better than both Ob and Vb. The Ob+Vb model effectively overcomes Ob's early and Vb's later limitations, in a manner similar to how the limitations observed in the constant speed and deceleration manoeuvers were addressed.

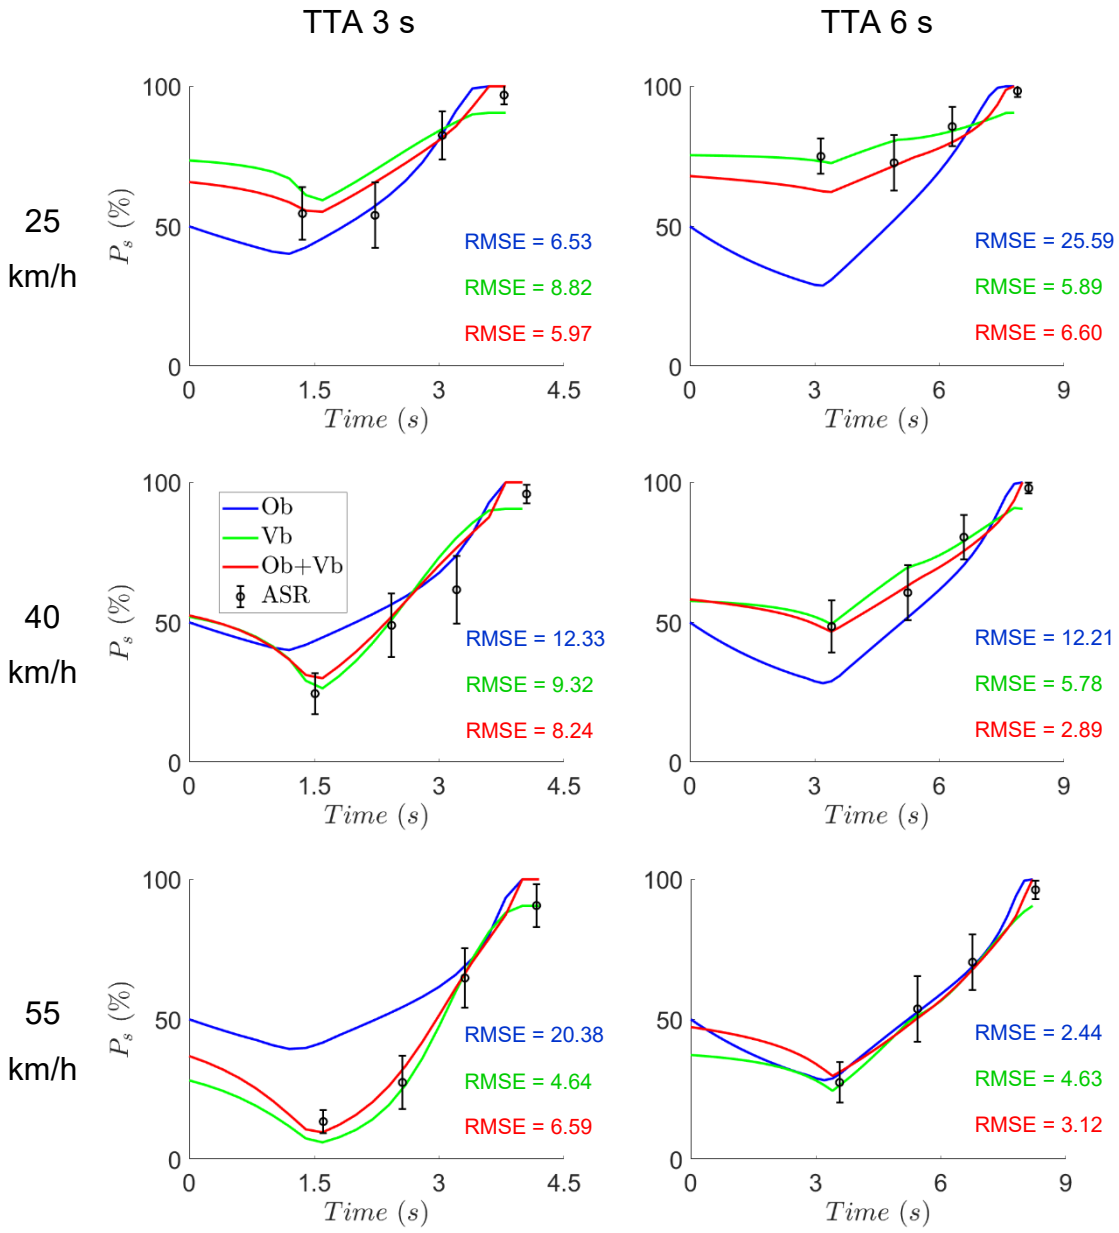

Figure 3.7 - Comparison between model predictions and pedestrians' beliefs for the mixed manoeuvres

Overall, the Ob+Vb model provided the highest Spearman's rank correlation (along with Vb), the lowest RMSE and, based on qualitative assessment supported by quantitative error values, appeared to be the most accurate in predicting the pedestrians' beliefs across all 18 kinematic scenarios. Even though there are strong indications that Ob+Vb is the best model (out of the three suggested ones), it is not possible to select it, before a deeper investigation (Subsection 3.4.2) and a model selection technique that prevents overfitting (Subsection 3.4.3) are presented. Overfitting occurs when a model fits the training data too closely, capturing noise and random fluctuations rather than the underlying patterns. This often happens when a model has too many free

parameters relative to the amount of data. In this case, Ob has 7 parameters, Vb has 5, and Ob+Vb has 12. The Ob+Vb model, with its significantly higher parameter count, carries a greater risk of overfitting compared to the simpler Ob and Vb models. Therefore, a model selection technique that penalises model complexity and assesses performance is necessary to ensure the selected model accurately reflects the underlying cognitive processes rather than being a mere reflection of the specific dataset. This will help us determine if the improved fit of the Ob+Vb model is due to genuine cognitive mechanisms or an artifact of its increased flexibility.

3.4.2 Breaking down the evidence

To better understand how the models are providing their predictions, this subsection delves deeper into the importance of the approaching vehicle's behaviour evidence. As a reminder, A_{ns} refers to non-stopping behaviour evidence, A_s refers to stopping behaviour evidence, A_{Ob} refers to behaviour evidence through the observation-based (Ob) approach and A_{Vb} refers to behaviour evidence through the value-based (Vb) approach (as an example A_{Vbs} refers to the value-based decelerating behaviour evidence). The following equations are provided for further clarity:

$$A_{ns} = A_{Obns} + A_{Vbns} \tag{3.15}$$

$$A_s = A_{Obs} + A_{Vbs} + M \tag{3.16}$$

Figure 3.8 presents the effects of different initial speeds and different vehicle manoeuvres on the vehicle behaviour evidence, for all scenarios with an initial TTA of 3 s. In scenarios with the same initial speed, the initial values of A_s, A_{ns} and $A_s - A_{ns}$ are the same. This is because: a) $A_{Obns} = A_{Obs} = 0$ since there are no prior observations leading towards any of the two behaviours, b) M is constant and c) because the initial speed, the required deceleration rate to stop before the pedestrian and the politeness constant are the same in all these cases, meaning that A_{Vbns} and A_{Vbs} also remain the same across these cases.

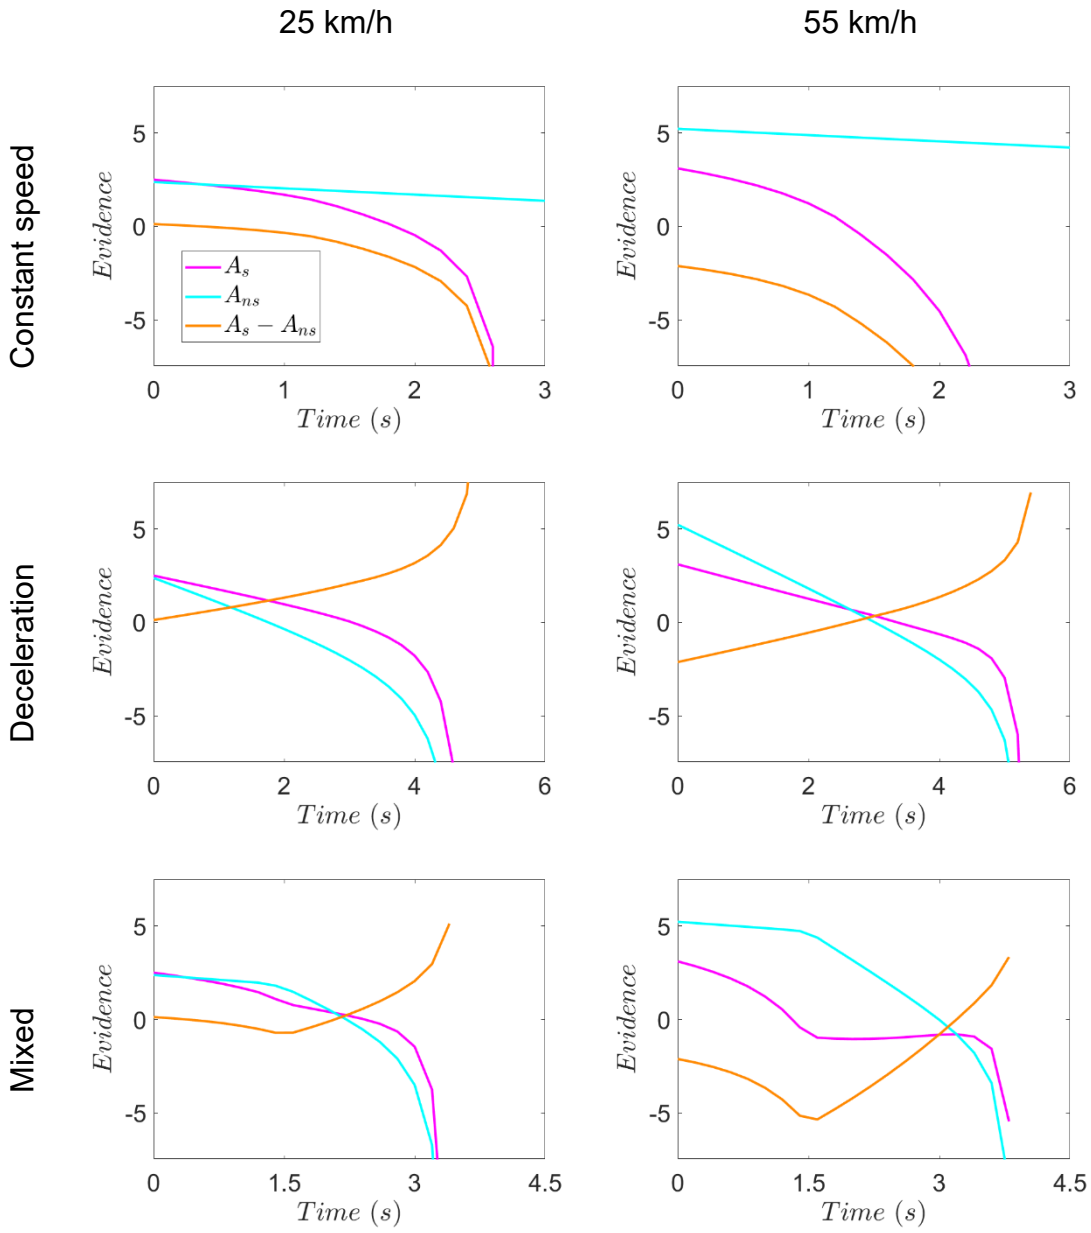

Figure 3.8 - Evidence of the two possible behaviours. The initial TTA is 3 s for all panels

Looking across scenarios with different initial speeds, the effect that has been discussed previously can be seen: When the initial vehicle speed is higher, there is a shift towards initially believing that the vehicle will maintain its speed rather than stop. This is visible in Figure 3.8 as $A_s - A_{ns}$ is initially more negative for higher initial speeds. This happens in the model because higher speeds imply a greater reward for the driver/vehicle if the vehicle maintains its speed, leading to a stronger initial bias towards non-yielding behaviour.

In a similar manner, Figure 3.9 illustrates the effects of different initial TTAs and different vehicle manoeuvres on the vehicle behaviour evidence, for all presented scenarios having an initial vehicle speed of 40 km/h. In scenarios with the same initial TTA, the initial values of A_{ns}, A_s and $A_s - A_{ns}$ are the same, for the same exact reasons as described above, regarding scenarios of the same initial speeds in Figure 3.8.

Looking across scenarios with different initial TTAs, another effect that has been discussed previously can be seen: When the vehicle is initially further away (longer initial TTA), there is a shift toward initially believing that the vehicle will stop rather than maintain a constant speed and not stop. This is visible in Figure 3.9 as $A_s - A_{ns}$ is initially less negative for longer initial TTAs. This happens in the model because larger initial TTAs, mean larger distance between the vehicle and the pedestrian, which means that the required deceleration rate for the vehicle to stop before the pedestrian's position is lower, thus the overall reward for the decelerating behaviour is larger, leading to a stronger initial bias towards the stopping behaviour.

However, the main takeaway from these two figures is that the $A_s - A_{ns}$ evidence difference is what dictates the shape of Ob+Vb's predicted beliefs (red curves in Figure 3.5, Figure 3.6 and Figure 3.7) regarding the approaching vehicle's stopping behaviour (P_s). The quantity $A_s - A_{ns}$, correlates directly with the behaviour probabilities predicted by the model. If $A_s - A_{ns} > 0$, then $P_s > P_{ns}$ and vice versa, meaning that the sign and magnitude of this difference indicate the relative belief towards the stopping versus the non-stopping behaviour.

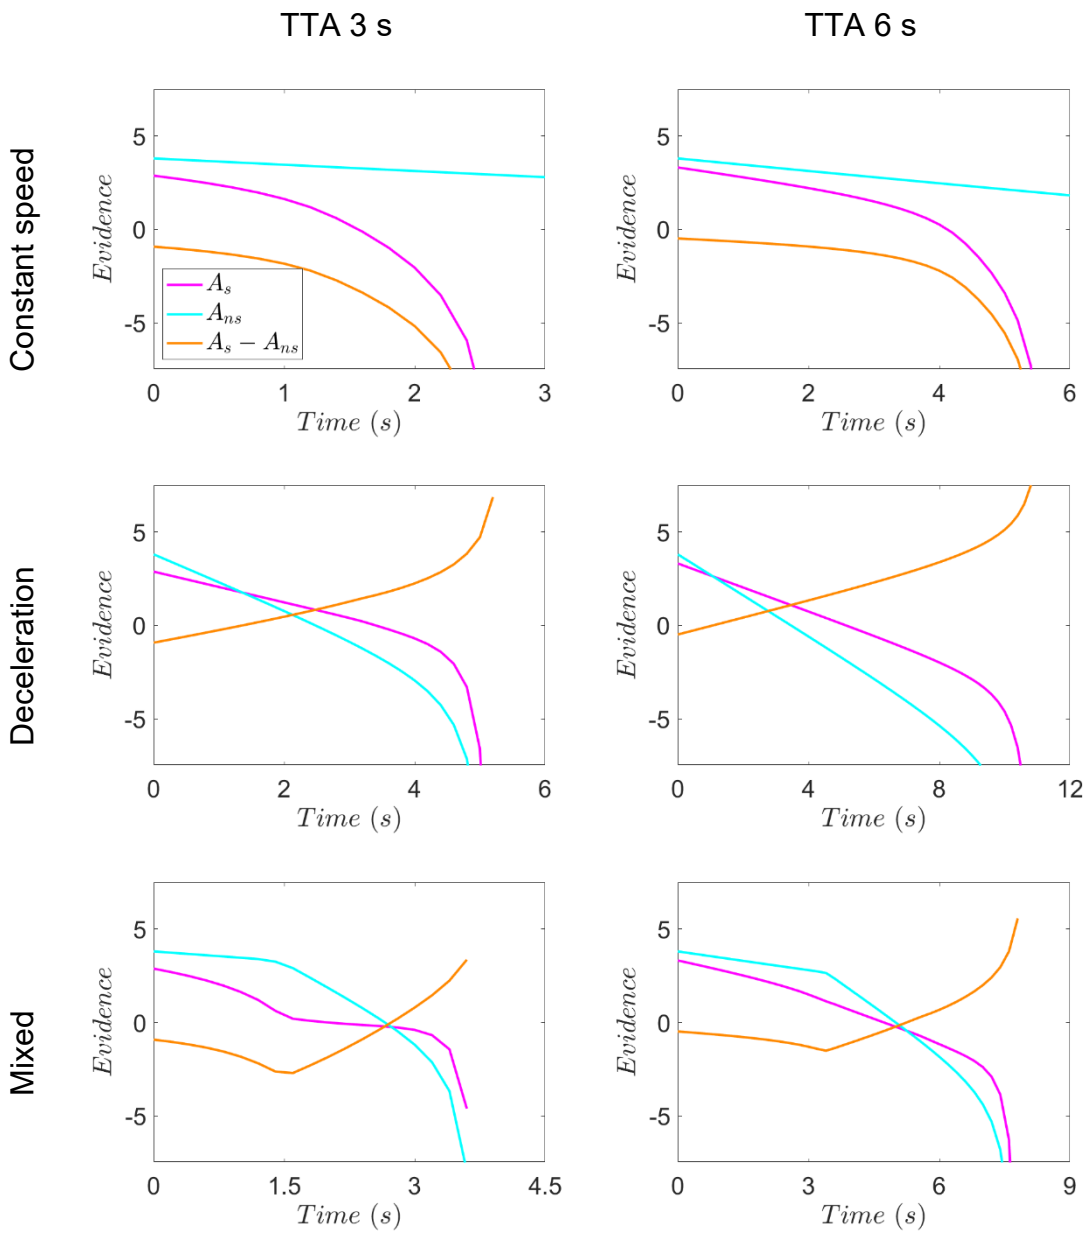

Figure 3.9 - Evidence of the two possible behaviours. The initial speed is 40 km/h for all panels

Building upon the analysis of vehicle behaviour evidence in the two previous figures, the individual contributions of the observation-based and value-based mechanisms are further explored in Figure 3.10 and Figure 3.11. To avoid any confusion the added soft-margin M has been omitted in Figure 3.10 and Figure 3.11.

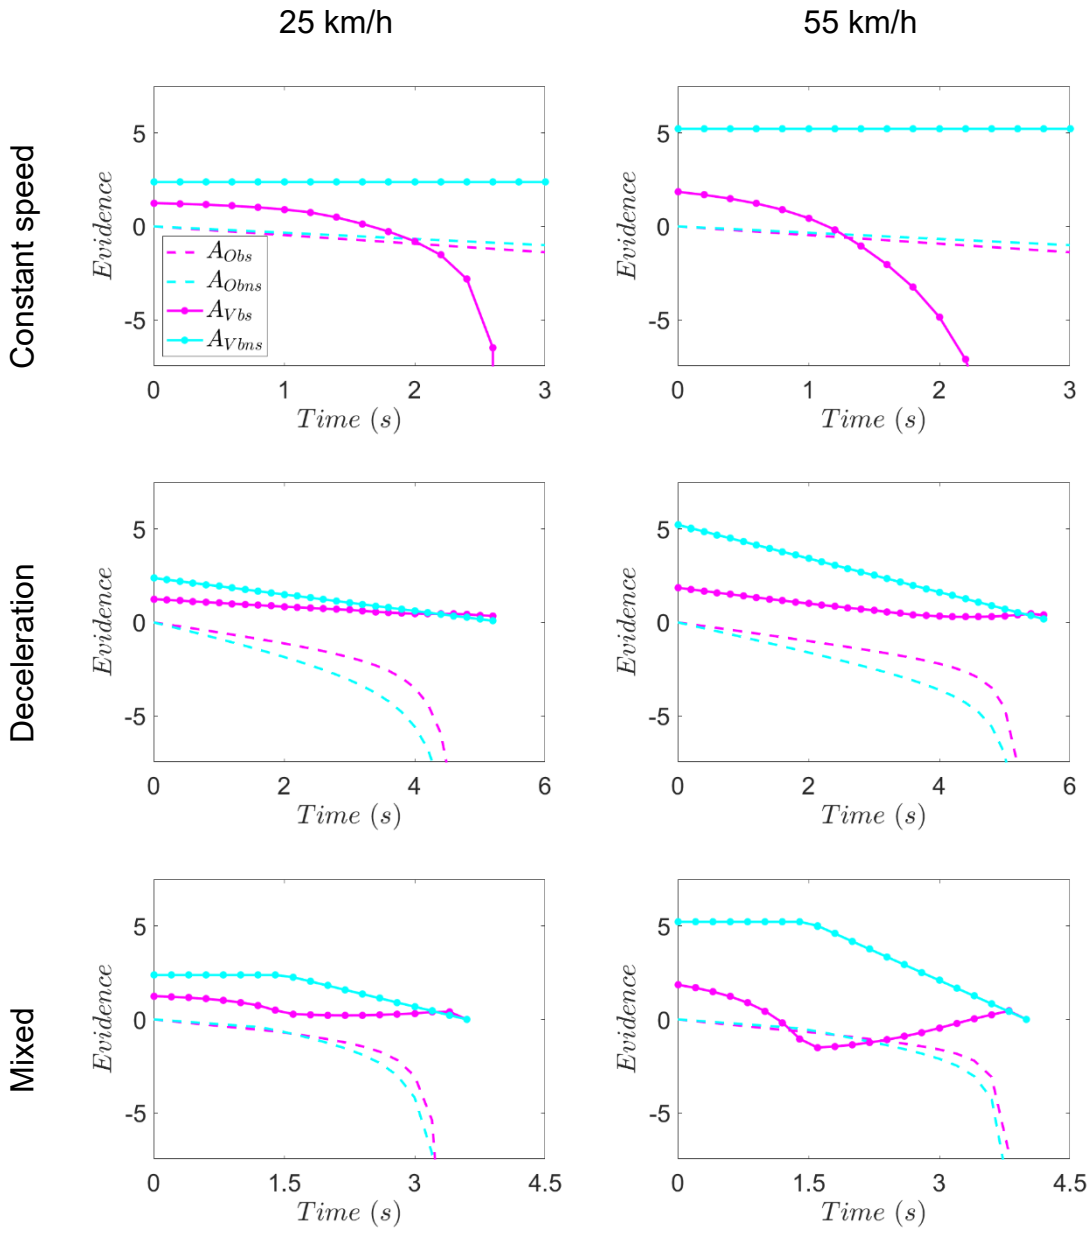

Figure 3.10 - Evidence of the two possible behaviours, divided by behaviour estimation mechanism. The initial TTA is 3 s for all panels

Specifically, Figure 3.10, which presents the same scenarios as Figure 3.8, illustrates the effects of different initial speeds and vehicle manoeuvers on the evidence calculated separately by the observation-based (A_{Obs} and A_{Obns}) and value-based (A_{Vbs} and A_{Vbns}) behaviour estimation approaches. Similarly, Figure 3.11, which mirrors the TTA variations seen in Figure 3.9, shows the effects of different initial TTAs and vehicle manoeuvers on the same evidence components.

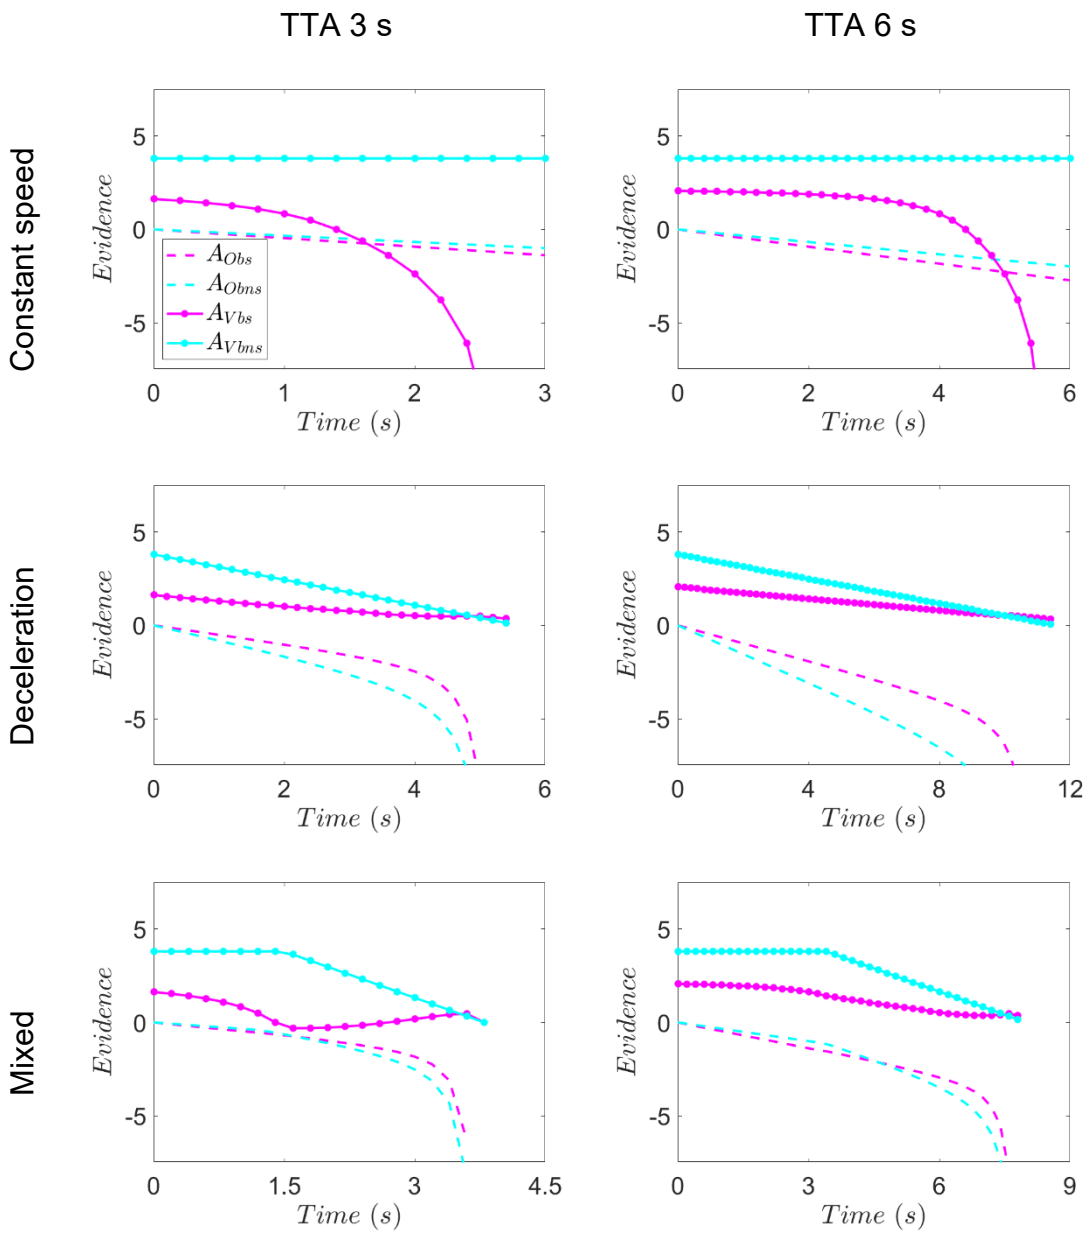

Figure 3.11 - Evidence of the two possible behaviours, divided by behaviour estimation mechanism. The initial speed is 40 km/h for all panels

Some key points can be drawn from Figure 3.10 and Figure 3.11. In the beginning of all scenarios, A_{Obs} and A_{Obns} both have zero evidence values since there is no available $\dot{\tau}$ observation, but A_{Vbs} and A_{Vbns} have non-zero evidence values. More specifically, $A_{Vbs}(t=0) > A_{Vbns}(t=0)$, which is understandable when considering Equations (3.4) and (3.5), as $R_s(t) = R_{ns}(t) - k_{da}a_{req}(t)^2 + c_{pol}$ (with c_{pol} not being substantial enough to compensate for the deceleration discomfort cost). This means that the estimated value for the car of keeping a constant speed is always initially greater than the estimated value of yielding. Since Ob's performance is based on A_{Obs} and A_{Obns}, it is apparent how this mechanism is limited in accurately capturing the initial beliefs of the pedestrians.

Conversely, Vb's performance is based on A_{Vbs} and A_{Vbns}, and has been seen that the effects of different speeds and TTAs on the initial average pedestrian beliefs are captured.

In the constant speed manoeuvres and the constant speed phase of the mixed scenarios (Figure 3.10 and Figure 3.11), A_{Vbns} remains constant since the vehicle's speed is constant, A_{Vbs} decreases towards $-\infty$ since the required deceleration to stop before the pedestrian increases towards $+\infty$ as the vehicle gets closer to the pedestrian's position. During the same periods of time, the observation-based evidence difference is $A_{Obns} - A_{Obs} > 0$ and increasing, as $\dot{t}$ observations keep being in line with the true vehicle behaviour (constant speed).

In the deceleration manoeuvres and the decelerating phase of the mixed scenarios (Figure 3.10 and Figure 3.11), A_{Vbns} decreases linearly towards 0 since the vehicle's speed decreases linearly towards zero. A_{Vbs} decreases towards the value of c_{pol} since the vehicle's speed and required deceleration decrease towards 0 (so in the end $A_{Vbns} - A_{Vbs} \gtrapprox 0$). During the same periods of time, the observation-based evidence difference is $A_{Obs} - A_{Obns} > 0$ and increasing (with very large values in the end), as $\dot{t}$ observations keep being in line with the true vehicle behaviour (stopping).

Looking at the evidence values in Figure 3.10 and Figure 3.11, it is possible to understand why Vb's performance is limited at the end of deceleration and mixed scenarios and why Ob performs better. Regarding Vb, as the vehicle decelerates, A_{Vbs} and A_{Vbns} converge, meaning that $A_{Vbs} \cong A_{Vbns}$, reducing the model's ability to capture the average pedestrian beliefs. In contrast, Ob's evidence values (A_{Obs} and A_{Obns}) continue to diverge, allowing it to accurately predict the final-stage beliefs, in the cases when the vehicle eventually stops before the pedestrian's position.

Therefore, this subsection confirms the benefit of combining Ob and Vb into Ob+Vb. The analysis has demonstrated that this combination is necessary by showing the importance of different sources of evidence and the distinct benefits of each behaviour estimation mechanism.

3.4.3 BSCV model selection

In order to select which model (amongst Ob, Vb, and Ob+Vb) had the best predictive performance, a Bootstrap Cross-Validation (BSCV) analysis was

performed, as described in Section 3.3. This analysis was deemed appropriate as it provides the ability to measure the goodness-of-fit of the models to data while it mitigates overfitting and controls model complexity – issues that are particularly relevant for the Ob+Vb model, which has a significantly larger number of parameters.

As mentioned in Section 3.3, the first BSCV analysis targeted Spearman correlations, fully in line with the Baker et al. (2009) approach. Figure 3.12 illustrates the distributions of the BSCV correlations for the three suggested models. From Figure 3.12, it is evident that Ob is the worst performing model. However, the distributions of Vb and Ob+Vb exhibit substantial overlap, making it impossible to definitively conclude which of the two correlates more strongly with the pedestrians' beliefs. For that reason, it was deemed necessary to perform a second BSCV analysis, but this time using RMSE as the goodness-of-fit measure (see Section 3.3).

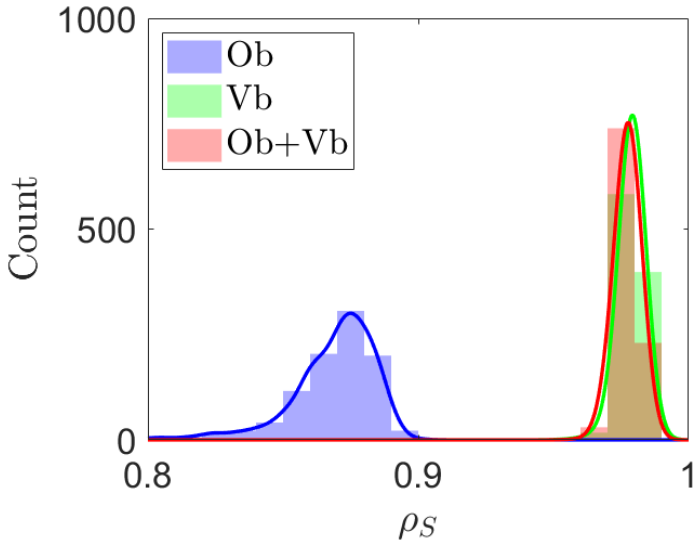

Figure 3.12 - Histograms of the BSCV correlations of Ob, Vb and Ob+Vb. The solid lines are the histograms' fitted distributions, which were obtained by MATLAB's Kernel smoothing function

In Figure 3.13, the distributions of the BSCV RMSEs for the three suggested models are presented. Once more, Ob is the worst of the three models, exhibiting the highest RMSE. When comparing Vb's and Ob+Vb's BSCV RMSEs, the picture becomes clearer in comparison to the BSCV Spearman correlations (Figure 3.12), with the Ob+Vb model showing a distribution shifted towards lower error values. This BSCV analysis aimed to assess the generalisation performance of the models by evaluating their goodness-of-fit on multiple

resampled datasets. To draw definitive conclusions about which model performs best, the distributions of the goodness-of-fit measures (in this case, RMSE) should ideally exhibit minimal or no overlap. Still, Figure 3.13 shows a noticeable overlap between the Vb and Ob+Vb distributions. Although the Ob+Vb distribution tends towards lower RMSE values, the overlap between the two distributions implies that the differences in model performance (minimisation of RMSE between predicted beliefs and average subjective ratings) between Vb and Ob+Vb might not be substantial or statistically significant.

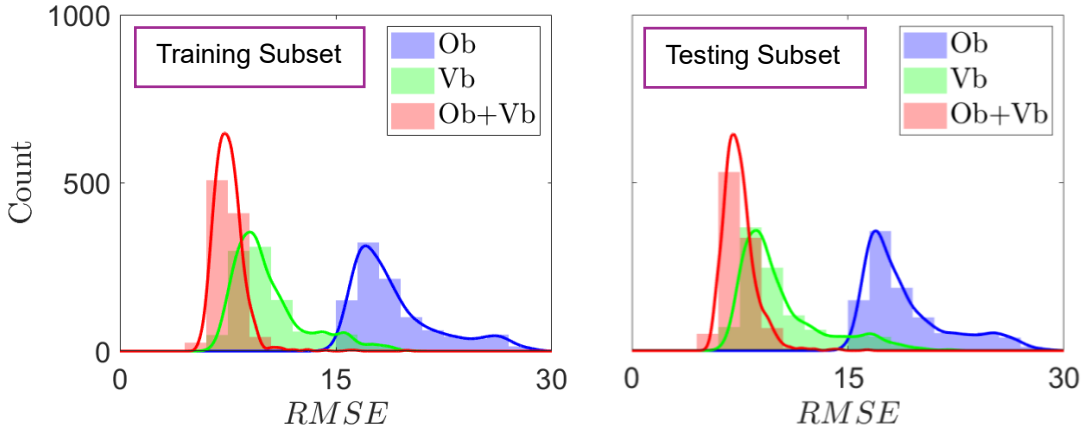

Figure 3.13 - Histograms (and their kernel-smoothed distributions) of the BSCV RMSEs of Ob, Vb and Ob+Vb

Therefore, based on the BSCV RMSE analysis, Ob is clearly the worst performing model, while Ob+Vb appears to provide a slightly better fit to unseen data compared to Vb, as it tends to produce lower prediction errors on average. The results are not fully conclusive because of the noted overlap between the Vb and Ob+Vb distributions. Additional tests would be needed to evaluate the significance of the observed differences and determine whether the better performance of the Ob+Vb model is meaningful and not due to overfitting.

3.4.4 Parameter investigation

To further validate the models and gain a deeper understanding of their behaviour, an in-depth parameter investigation was performed. This analysis examined the range of parameter values that yielded good model fits during the Bootstrap Cross-Validation (BSCV) process. Unlike the grid search method that was applied in Subsection 2.5.3, through which an extensive range of parameter settings was investigated, the current analysis was based on the parameter values obtained directly from the BSCV analysis.

As explained in Section 3.3, the BSCV technique fitted the suggested models to resampled datasets. Applying the BSCV model selection, generated a distribution of values for every parameter of the tested model. That distribution represented the range of parameter values under which the model performed best in fitting on the resampled data. These distributions, which are bootstrap estimates, indicate the likely values of the parameters that would generalise well to unseen data. By examining these parameter distributions, the robustness of the models' parameters can be assessed and the sensitivity of the model predictions to variations in parameter values can be further investigated.

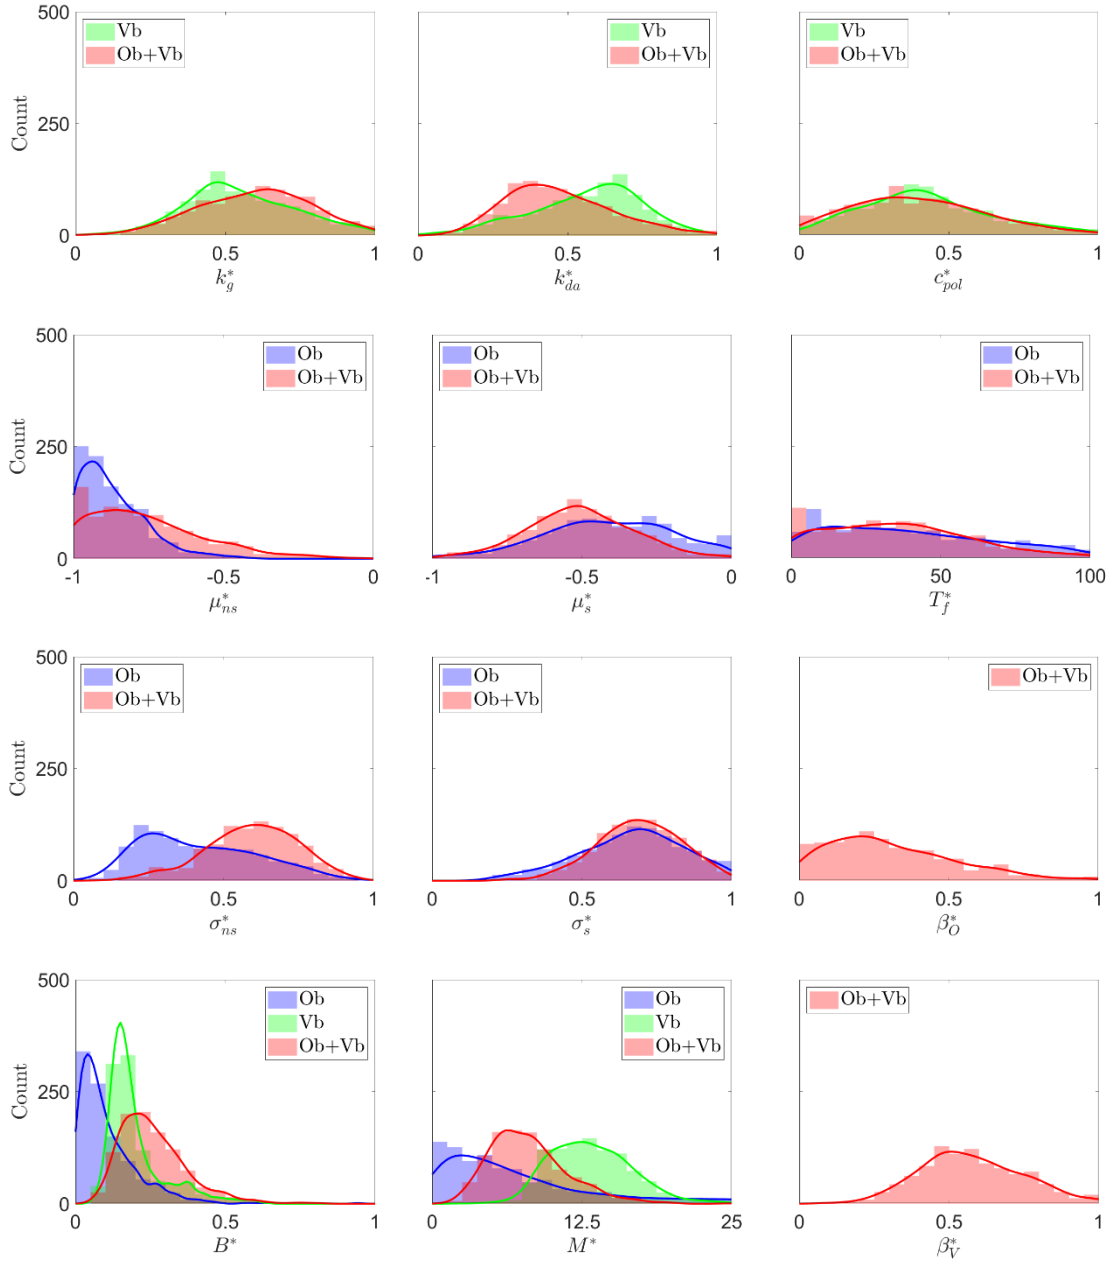

Figure 3.14 - Histograms (and their kernel-smoothed distributions) of the BSCV-obtained parameters of Ob, Vb and Ob+Vb

Figure 3.14 shows the wide range of parameter values, close to the best-fitting values, which yielded the lowest possible BSCV RMSEs for each model. Ob best captured pedestrians' beliefs at low values of μ_{ns} ($\dot{\tau} = -1$ indicates constant speed), intermediate values of μ_s ($\dot{\tau} = -0.5$ indicates adequate deceleration rate to stop the vehicle safely before a pedestrian, avoiding a collision), low values of σ_{ns} (narrow constant speed behaviour distribution), intermediate values of σ_s (wider decelerating behaviour distribution), all values of T_f (with denser area in lower values – not forgetting past observed evidence quickly), low values of B (softer softmax evidence transformation, meaning less sensitivity to differences in behaviour evidence) and low values of M (almost no prior belief towards the decelerating behaviour). It is worth noting that M for Ob was fitted across both positive and negative values. However, the optimisation algorithm consistently found that, in both cases, M values close to 0 yielded the best Ob fits. This strongly suggests that the Ob model, as parameterised, does not incorporate any prior belief shift towards either decelerating or constant speed behaviour. The algorithm's consistent selection of $M \cong 0$ reinforces the idea that, within Ob, the initial beliefs are purely driven by the observational $\dot{\tau}$ information, not by any inherent bias towards one behaviour over the other.

Vb captured pedestrians' beliefs best at parameter values that highlight the relative importance of the following factors: the approaching vehicle's speed (k_g), the required deceleration for the vehicle to stop before the pedestrian (k_{da}), and the vehicle exhibiting pro-social behaviour (c_{pol}). The model also favoured a relatively softer softmax evidence transformation (B) and a substantial prior belief towards decelerating behaviour (M).

Ob+Vb captured pedestrians' beliefs best at parameter values that revealed a nuanced integration of observational and value-based information. Comparing these values to those of Vb and Ob offers some key insights. Firstly, Ob+Vb exhibited a slightly higher importance of the approaching vehicle's speed (k_g) and a slightly lower importance of the required deceleration (k_{da}) compared to Vb. This suggests that when both observation and value-based information are available, the model places a greater emphasis on the immediate speed of the vehicle as a predictor of its behaviour, potentially because speed provides a more immediate and salient cue. Secondly, Ob+Vb, like Ob, favoured lower values of

μ_{ns} and intermediate values of μ_s, aligning with the theoretical expectations of the observation information ($\dot{\tau}$). Thirdly, Ob+Vb showed the same rate of forgetting past observed evidence as Ob (T_f), indicating a similar reliance on recent observations between these two models. The parameters β_O (equal to 0 would mean purely value-based behaviour estimation) and β_V (equal to 0 would mean purely observation-based behaviour estimation) act as weights, balancing the contributions of the value-based and observation-based evidence. Due to their role as relative scaling factors, one of these two parameters could be fixed without altering the model's fundamental behaviour. For instance, fixing β_O to 1 would allow β_V to represent the relative weight of the value-based evidence compared to the observation-based evidence. This approach would simplify the model by reducing the number of free parameters and potentially improve the interpretability of the remaining parameters. However, it's important to acknowledge that fixing a beta parameter would likely affect the optimal values of the other parameters. In this analysis, none of the beta parameters were fixed. Finally, Ob+Vb favoured a slightly harsher softmax evidence transformation (B) than both Ob and Vb. This implies that Ob+Vb is more sensitive to small differences in evidence, leading to sharper transitions in beliefs. This is consistent with the model's integration of both observation and value-based information, which might lead to more confident and decisive belief updates.

Figure 3.15, Figure 3.16 and Figure 3.17 illustrate the relationships between all the pairs of parameters and how they affect the overall fit of the model. Based on the Bootstrap Cross-Validation (BSCV) analysis, these figures show how similar model fits (measured by RMSE) can be obtained by a wide range of BSCV parameter combinations for each of the three models (Ob, Vb, and Ob+Vb). Specifically, these figures display the same information as the single-parameter distributions shown in Figure 3.14, but now across two parameter dimensions at a time. The diagonal panels of Figure 3.15, Figure 3.16 and Figure 3.17 illustrate the single-parameter distributions, which are identical to those seen in Figure 3.14. The pairwise scatterplots (off-diagonal panels) present the relationships between the respective pairs of parameters.

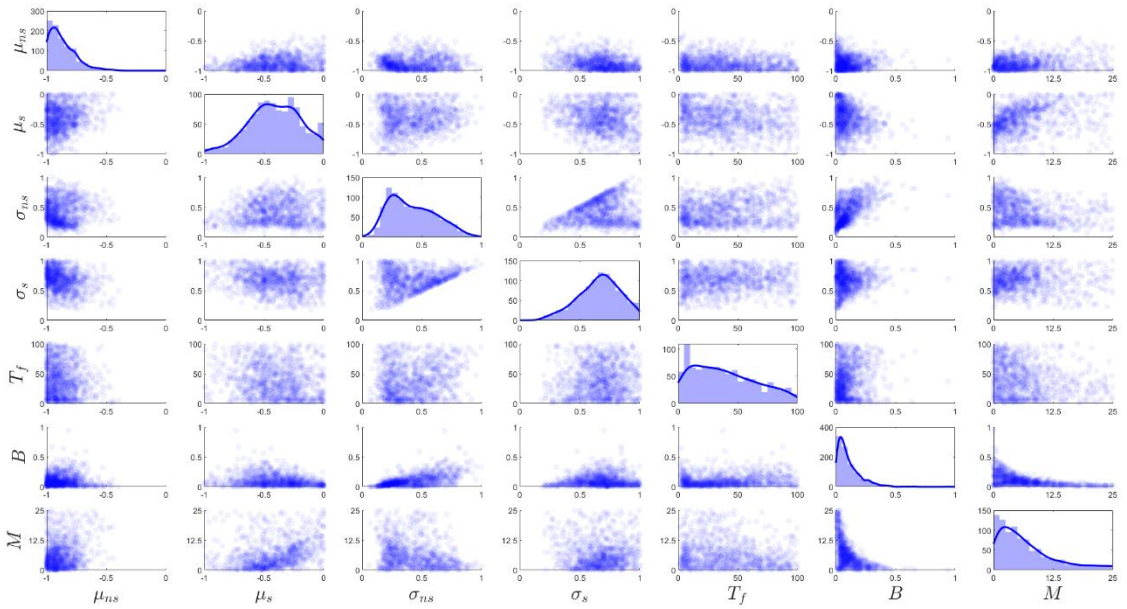

Figure 3.15 - Ob's pairwise parameter scatterplot matrix. The histograms in the diagonal illustrate the distribution of the respective parameter values with the BSCV-obtained RMSEs. The scatterplots show the pairwise parameter combination areas with obtained RMSEs

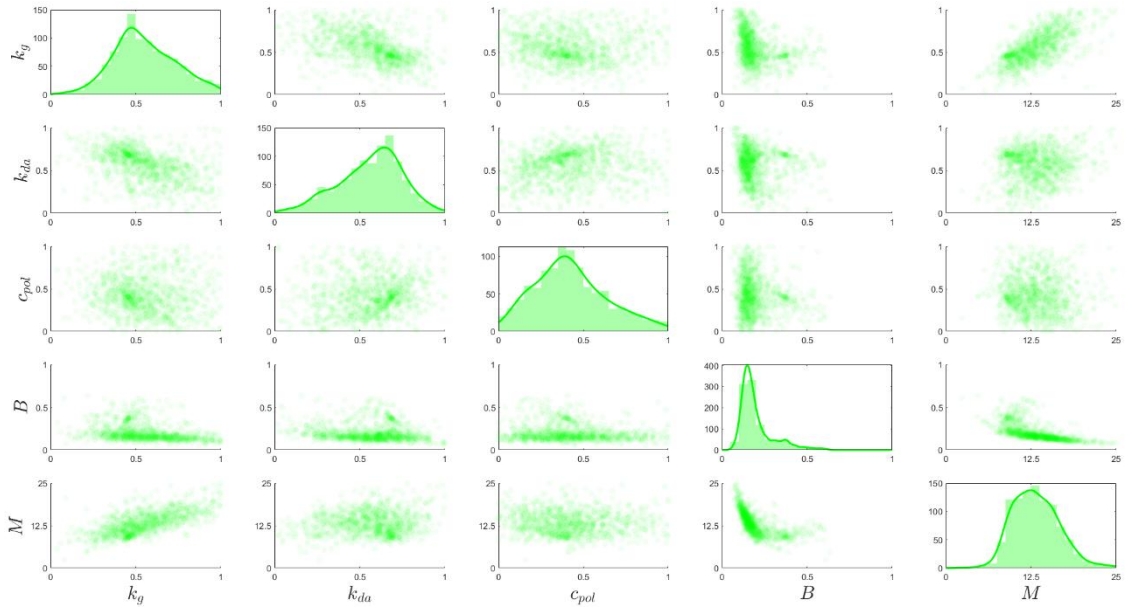

Figure 3.16 - Vb's pairwise parameter scatterplot matrix. The histograms in the diagonal illustrate the distribution of the respective parameter values with the BSCV-obtained RMSEs. The scatterplots show the pairwise parameter combination areas with obtained RMSEs

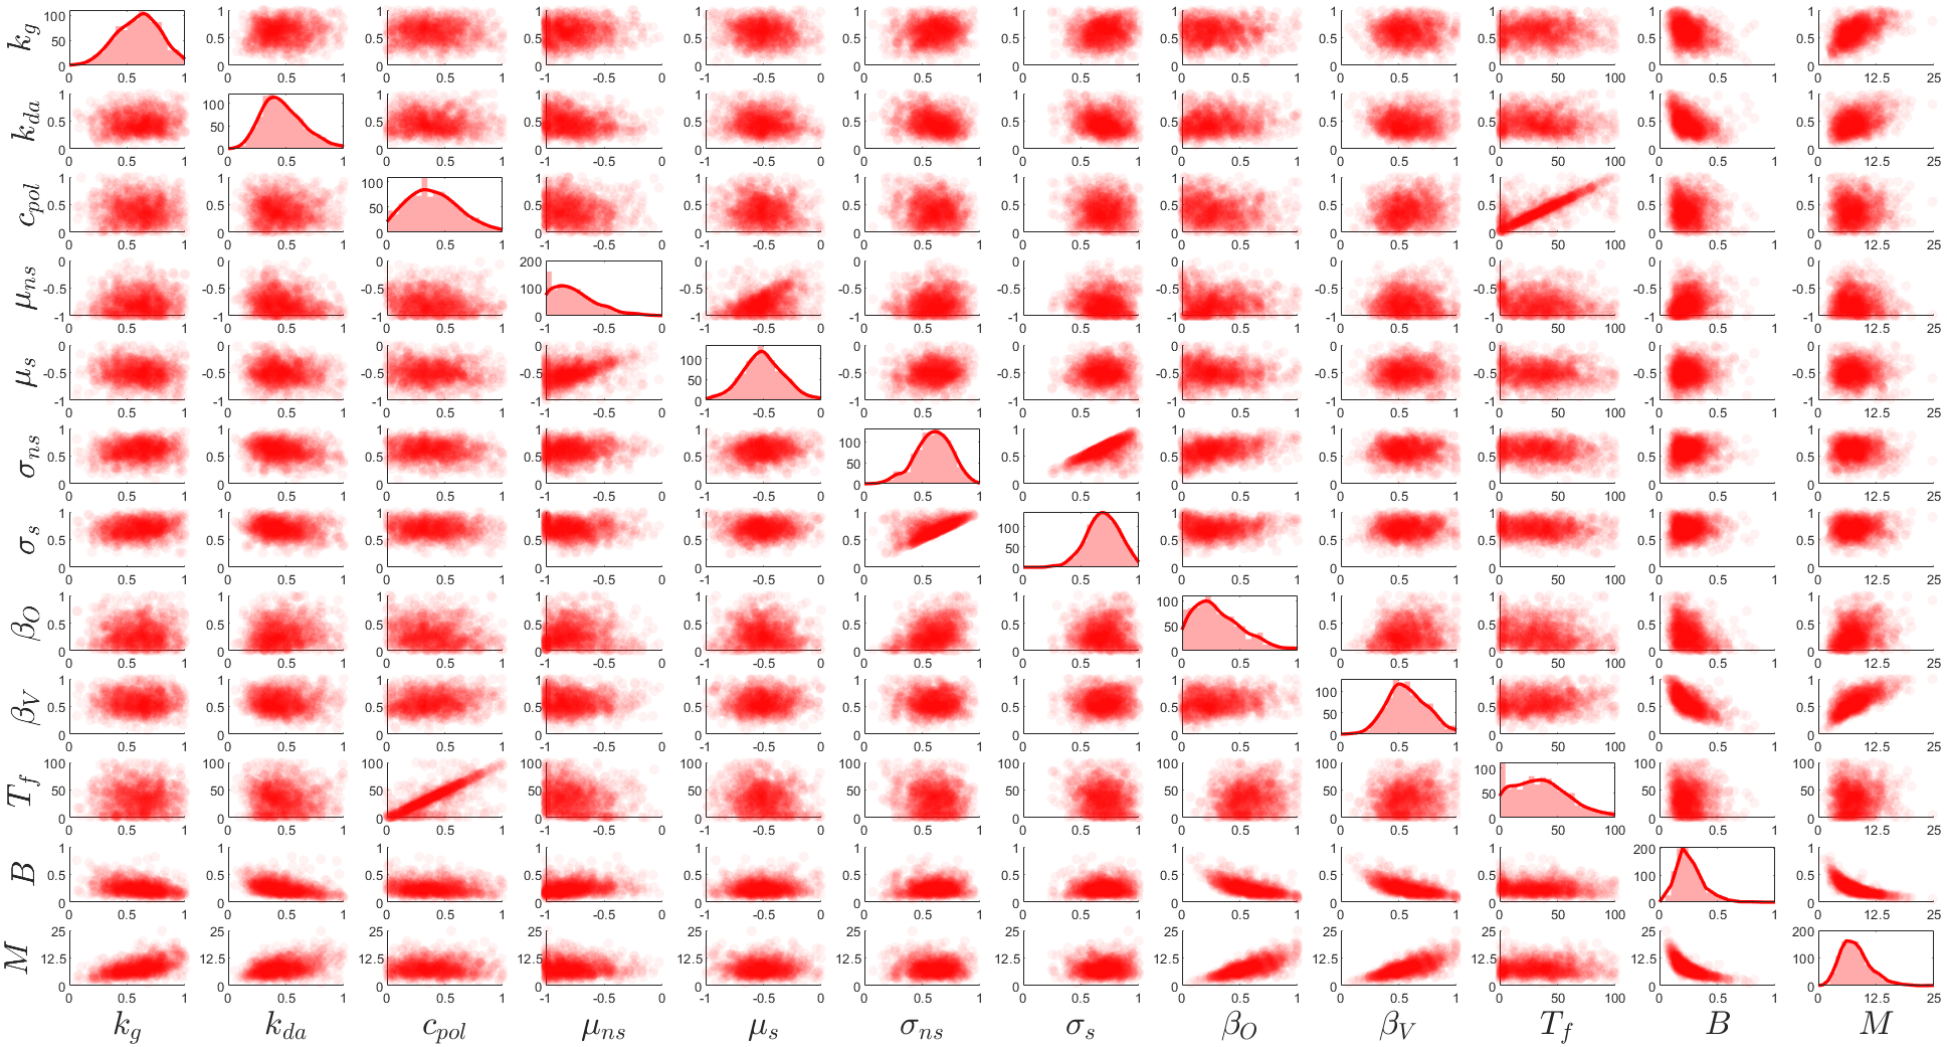

Figure 3.17 - Ob+Vb's pairwise parameter scatterplot matrix. The histograms in the diagonal illustrate the distribution of the respective parameter values with the BSCV-obtained RMSEs. The scatterplots show the pairwise parameter combination areas with obtained RMSEs

The fact that multiple parameter combinations yield roughly equivalent model fits indicates that the parameterisation is likely redundant. This means that the model's performance is not highly sensitive to the precise values of individual parameters, but rather to the overall balance between them. In other words, several sets of parameter values can produce similar predictions. This redundancy limits the interpretability of the frequency distributions (the single-parameter distributions along the diagonal) and the pairwise scatter plots (the off-diagonal panels). Because of the observed parameter correlations, interpreting single-parameter distributions in isolation can be misleading. For instance, if two parameters are highly correlated, a change in one parameter might necessitate a corresponding change in the other to maintain optimal model fit. Figure 3.15, Figure 3.16 and Figure 3.17 indicate correlations between some parameters.

Understanding the behaviour of the models depends much on the different levels of redundancy found by the parameter analysis among Ob, Vb, and Ob+Vb.

In the Ob model, a positive correlation was observed between σ_s and σ_{ns}, indicating that the model can achieve similar fits with proportional adjustments to these parameters. Additionally, a negative correlation between T_f and B could possibly suggest that the duration of forgetting observation-based evidence and the internal translation of behaviour evidence into beliefs might be changed in a compensatory way. The wide distributions of these parameters reinforced the idea that the model's performance is not highly sensitive to their precise values.

The Vb model displayed a different pattern of redundancy. B and M showed a negative correlation between them, indicating that Vb may balance the internal translation of behaviour evidence into beliefs and the prior bias towards the stopping behaviour. Vb also demonstrated flexibility in the weights assigned to maintaining constant speed to reach a goal and the deceleration discomfort. However, the rather weak correlations between the other parameters, suggest that the performance of the model might be more sensitive to individual parameter values in these cases.

The Ob+Vb model displayed the highest degree of parameter redundancy, with numerous correlations observed between various parameter pairs. Similar to the Ob model, a strong positive correlation between σ_s and σ_{ns} was observed,

along with a potential negative correlation between T_f and B. However, the Ob+Vb model also exhibited correlations between β_O, β_V and other parameters, particularly with T_f, B and M. This suggests that the weighing of observation-based and value-based evidence is intricately linked with the evidence forgetting window, the internal translation of behaviour evidence into beliefs and prior beliefs. It is acknowledged that the model may be overspecified with respect to the evidence weight parameters. While reducing the parameter space by fixing one of them (e.g., $\beta_V = 0.5$) would ensure better identifiability, this approach was avoided in favour of the full theoretical expressiveness of their respective roles in the evidence summation function. Furthermore, while correlation between other parameter pairs was observed, Figure 3.17 did not show any clear correlation between β_O and β_V. Therefore, both evidence weights were retained.

The wide parameter ranges observed for most parameters in all three models indicate that the models' performance can remain unaffected by a variety of different parameter values. However, the complexity of compensatory adjustments varies across models. The Ob and Vb models exhibit relatively simpler forms of redundancy, primarily involving two parameters at a time. In contrast, the Ob+Vb model displays a more intricate network of correlations, suggesting that the model's behaviour is driven by the overall balance between multiple parameters.

The high degree of redundancy in the Ob+Vb model poses a significant challenge for interpreting individual parameters. The model's flexibility and complexity make it important to focus on the overall patterns and relationships between parameters rather than individual values. For that reason, emphasis must be put on the interplay between different sources of information and cognitive factors when interpreting the model's predictions.

In summary, while all three models exhibit parameter redundancy, the Ob+Vb model stands out for its complex connection of correlations and flexibility, reflecting its simultaneous integration of both observation and value-based evidence. This increased complexity necessitates a cautious approach to parameter interpretation, emphasising the importance of considering the model's overall behaviour while also requiring further investigation of its ability to generalise to previously unseen data and scenarios.

3.5 Discussion

The work presented in this chapter showed that the combination of observation-based and value-based behaviour estimations can predict the average pedestrian beliefs quite accurately, across all 18 tested kinematic scenarios. Firstly, a Bayesian observer model, formulated and tested by Baker et al. (2009), was adapted to the traffic scenario presented in Chapter 2. In Chapter 3, that model was modified in order to calculate observation-based evidence for the two possible vehicle behaviours (decelerating to stop and maintaining constant speed with no intention to stop). In the current chapter a value-based behaviour estimation mechanism was implemented, based on the pedestrians' expectations that the driver of the approaching vehicle (or AV) would try to maximise the value/reward of their own behaviour, originally proposed by Markkula et al. (2023). Lastly, the two behaviour estimation mechanisms were combined into a more complex model. Since the number of parameters increased substantially, new model fitting and validation methods were needed. To address these two needs, a global controlled random search + local line search optimisation algorithm and the BSCV technique were adopted and applied.

The results in this chapter have shown that, qualitatively speaking, all three models were able to capture the majority of the patterns of the pedestrians' beliefs. More specifically, Ob had the worst predictive performance out of the three, which may be because it cannot account for any kinematics-related effects other than deceleration ($\dot{\tau}$ in this particular case). Since this behaviour estimation model has no access to speed and TTA, its predictions can only change when $\dot{\tau}$ is changing, which is not true for the human data. This means that Ob provided poor predictions (the same predicted belief curve) in the constant speed scenarios and constant speed phases of the mixed scenarios. On the other hand, Vb was able to accommodate for what Ob was missing. With the use of rewards related to speed (vehicle's progress towards the goal), TTA/distance and deceleration (deceleration discomfort) and pro-social behaviour (yielding politeness), the majority of kinematic conditions' effects and belief patterns were captured by Vb. However, Vb seemed to be limited in predicting the pedestrians' beliefs at the end of vehicle approaches, when the vehicle was coming to a full stop. It was shown that the Vb model struggled at predicting pedestrians' beliefs

when the vehicle was travelling with relatively low speed because in that type of situation, the value-based evidence for the stopping and non-stopping behaviours are approximately equal, as the difference between these sources of evidence is almost equal to zero. Looking back at the importance of behaviour evidence (Subsection 3.4.2) for forming and updating the pedestrian beliefs, $A_s - A_{ns} = 0$ creates an uncertainty for the model, when in reality the stopping behaviour of the vehicle is quite apparent. Ob+Vb capitalised on the strengths of the two behaviour estimation mechanisms mentioned above. It was able to capture all the kinematic conditions' effects and belief patterns which were present in the current dataset.

Quantitatively, Ob+Vb and Vb achieved nearly perfect positive Spearman's rank correlations, while Ob's correlation was lower but still relatively high. In terms of prediction accuracy, Ob+Vb achieved the lowest RMSE, followed by Vb and then Ob. These findings were somewhat supported by the BSCV analysis, which again showed that Ob+Vb and Vb obtained the highest BSCV correlations and overall lower levels of RMSEs. However, the picture of the BSCV RMSEs was not very clear in order to decide whether Vb's or Ob+Vb's predictive performance was better, since their BSCV RMSE distributions overlapped, despite Ob+Vb tending to illustrate lower BSCV RMSEs than Vb on average. These findings suggest a need for further model comparison to be able to select the model with the best predictive performance.

However, Ob+Vb's possible high complexity, also indicated by its observed parameter redundancy, warrants further discussion. The pairwise correlations between its parameters suggest that the model's parametrisation is likely redundant and its performance likely relies on an overall balance between the values of the multiple interacting parameters, making individual parameter interpretation challenging. Although the model's complexity yields higher predictive accuracy, it demands a thorough analytical approach, as parameter redundancy could hinder its generalisability. To address this uncertainty, in the next chapter, a new experiment with new untested scenarios was designed, to validate the findings of Chapters 2 and 3, and test the models (most importantly Ob+Vb's) ability to generalise to unseen data.

4 Validating and generalising the behaviour estimation models

Building upon the model formulations and analyses presented in Chapter 3, this chapter focuses on validation and extrapolation of the developed behaviour estimation models. While in Chapter 3 the theoretical framework was established, the evaluation and the initial comparison of the observation-based (Ob), value-based (Vb), and augmented (Ob+Vb) models, in Chapter 4 the models' predictive capabilities were rigorously tested on a novel dataset, thereby assessing their generalisability.

To this end, a new experiment was designed, drawing upon the methodological principles established in Chapter 2, but incorporating key modifications to explore the models' performance under varied conditions. These varied conditions included scenarios directly comparable to those used in the experiment described in Chapter 2, allowing for the validation of the previous findings, as well as entirely new scenarios intended to test the models' predictive capabilities and assess their ability to extrapolate to other situations. This two-pronged approach intended to enable both a confirmation of previous results and an exploration of the models' behaviour in new untested scenarios. The data acquired from this experiment, in a similar manner to the previous experiment and analysis of Chapter 3, served as the basis for evaluating the predictive accuracy and robustness of Ob, Vb, and Ob+Vb. Consistent with the approach taken in Chapter 3, Bootstrap Cross-Validation (Baker et al., 2009) was employed to assess the goodness-of-fit of each model to the new dataset. This chapter details the design and implementation of the new experiment, presents the results of the model evaluations on the collected dataset, and discusses the implications of these findings for understanding and predicting pedestrian beliefs about approaching vehicle behaviour, ultimately validating and extending the findings presented in the previous chapters.

4.1 Experiment

As discussed in Chapter 2, the initial experiment was based on artificial kinematics that were focused on the idea: the vehicle is stopping or not stopping. Building upon this foundation and aiming to explore a wider range of realistic driver behaviours, this chapter introduces four manoeuvres grounded in both empirical observations, theoretical considerations and informal static driving simulator tests. This shift towards more varied and realistic manoeuvres allowed for a more nuanced understanding of pedestrian belief formation and updating. More details regarding the participants, the apparatus, the design and the procedure of the experiment are presented in the following subsections.

4.1.1 Participants

The study received ethical approval from the University of Leeds Ethics Committee (via an amendment request to the previous ethical approval, with reference LTTRAN-145). 30 healthy adults, including 14 males and 16 females, aged between 22 and 65 (mean age = 36.56, standard deviation = 10.68) were recruited from the University of Leeds Virtuocity participant list and Microsoft Teams channels. While the average age was slightly higher in this experiment than in the previous one, the participant selection criteria established in the experiment described in Chapter 2, this study maintained all the key requirements to ensure consistency; participants were required to have no significant mobility issues or medical conditions such as epilepsy, to have either normal or corrected-to-normal vision and have lived in the UK within the last 12 months, as their experience with road traffic could influence their road crossing behaviour and judgments. they provided written informed consent before participating and were given £15 as a reward for their participation, as in the first experiment.

4.1.2 Apparatus

The experimental apparatus, including the projection and tracking equipment, and virtual environment development software were identical to those described in Chapter 2, employing the HIKER lab's immersive, multi-projection virtual reality setup. Participants experienced the new simulated pedestrian scenarios within this environment.

4.1.3 Experimental design

The experimental design maintained the core principle of presenting participants with simulated vehicle approaches and requiring them to infer the vehicle's behaviour (stopping or not stopping) which was an adaptation of the Baker et al. (2009) original experiment in the pedestrian-vehicle interaction setting, by once again incorporating elements from the experiments of Pekkanen et al. (2022) and Dey et al. (2021), in a similar manner to Chapter 2. As stated in the introduction of this section, the main difference between the new experiment and the previous one, was that in the new experiment the tested driving manoeuvres were more varied and could be described as more realistic.

The simulated traffic environment, consistent with the previous experiment, featured a residential block with a 4.2 m wide, one-lane road and an unsignalised intersection during daylight hours. An autonomous blue sedan vehicle, travelled in the centre of the road, simulating the kinematics of the experimental design. The new experiment had the same independent variables categories as in the previous experiment: vehicle driving manoeuvres, initial Time-To-Arrival (TTA), and initial vehicle speed.

However, this experiment differed from the one described in Chapter 2 in the specific values used for initial vehicle speeds and types of driving manoeuvres. Specifically, the vehicle approached the pedestrian at initial speeds of 20 km/h, 40 km/h, and 60 km/h, instead of the 25 km/h, 40 km/h, and 55 km/h used previously. This adjustment was made for two reasons. One was to test the generalisability of the models and so it was decided to explore a broader range of speeds, by extrapolating on the low and high speeds, having in mind that these speeds can also be observed in real life interactions. The second was to validate the models and so it was decided to keep the intermediate (and quite frequent in real world interactions) speed the same as in the previous experiment. The initial TTAs were exactly the same as in the previous experiment (3 and 6 s). The

decision to retain TTA values of 3 and 6 s in the second experiment was based on their significance in pedestrian road crossing decision-making literature and the methodological constraints of the experimental paradigm. As stated before, this TTA range is recognised as a critical window in pedestrian gap acceptance. Also, using TTA values which are shorter than 3 s would have been problematic for the segmentation method. The process of truncating the scenario into four segments would lead to initial segments that are too short in duration to allow the participant to perceive the visual stimulus or for the segment kinematics to be meaningfully differentiated. Lastly, TTAs longer than 6 s were not of interest as pedestrians would almost always cross the road immediately without probably needing to interpret the vehicle's behaviour.

Furthermore, the vehicle's driving manoeuvres were also expanded. The first two of these manoeuvres were repeated from the first experiment in the light of validating the previous experiment and findings, while the second two manoeuvres were designed to test the models' generalisability to more complex and realistic driving situations. The design of the second two manoeuvres was influenced by the informal static driving simulator tests. These informal tests involved driving in a simulated urban environment. Drivers were reaching and maintaining a target speed (20, 40, or 60 km/h) and as soon as they were at a predetermined TTA (3s or 6s) to static virtual pedestrian (waiting at a bus stop), a beep informed them to perform specific braking manoeuvres. The obtained qualitative speed profiles (averaging repeated trials with the same kinematic conditions) helped define the kinematic details of these two manoeuvres, for which there is a lack of information in the literature. The four driving manoeuvres of this experiment are listed below and are illustrated in Figure 4.1.

- Asserting priority: Constant Speed Manoeuvre

This manoeuvre reflects a driver's tendency to prioritise their own progress, even when a pedestrian is present. Research indicates that a significant proportion of drivers maintain or even increase their speed when approaching a crossing with a competing pedestrian (Várhelyi, 1998). This behaviour can force pedestrians to yield, potentially leading to dangerous situations (Rasouli et al., 2018). This assertive approach contrasts sharply with yielding behaviours and highlights the potential conflict between driver expediency and pedestrian safety. This manoeuvre will be referred to as Constant Speed and is similar to the

constant speed scenarios from Chapter 2. In the constant speed scenarios, the vehicle maintained a consistent speed throughout the simulation.

- Yielding acceptance based on 1-stage braking: Constant Deceleration Manoeuvre (AV-like)

This manoeuvre, inspired by AV behaviour, emphasises predictability and smooth deceleration. It aligns with the concept of a "coordination smoother," where the vehicle's actions are designed to minimise variability and enhance predictability (Domeyer et al., 2020). This could be achieved through a consistent, linear deceleration towards a stop, signalling a clear intention to yield. Researchers have indicated that a vehicle's yielding behaviour (and thus its intention to stop) is recognised earlier when cues like lower speeds and higher deceleration rates are observed (Ackermann et al., 2019; Tian et al., 2023). So, with the constant deceleration manoeuvre the aim was to provide a clear and unambiguous signal to pedestrians. This manoeuvre will be referred to as Deceleration and is similar to the deceleration scenarios from Chapter 2. In the deceleration scenarios, the car decreased its speed at a constant rate until it stopped 2 to 2.5 meters from the participant.

- Yielding acceptance based on 2-stage braking: Mixed Deceleration Manoeuvre (Human-like)

This manoeuvre mimics human driver behaviour characterised by a two-stage deceleration process. Informal tests in a static driving simulator, along with previous findings (Lee et al., in prep), suggest that drivers often initially reduce speed by lifting off the gas pedal or applying a light deceleration, followed by a firmer deceleration to yield before reaching the pedestrian's position. This mixed approach reflects the complex sensorimotor communication between driver and pedestrian (Domeyer et al., 2020), potentially conveying a more nuanced intention to yield. This manoeuvre was informed by the informal static driving simulator tests and will be referred to as 'Two-stage Deceleration' and even though its kinematic details are different from the mixed scenarios of the previous experiment (Chapter 2), their general concept is similar. In the two-stage deceleration manoeuvre scenarios, the vehicle combined an initial phase of very subtle deceleration (almost constant speed, which was different from the purely constant speed phase of the mixed manoeuvres of the previous experiment by a

slight, controlled speed reduction) and then a second deceleration phase (with higher deceleration rates), to explore the impact of the vehicle's behaviour change, during the simulation, on the pedestrian's belief updating. In this scenario, the vehicle applied the subtle deceleration for a certain time (1.2 s for the 3 s initial TTA condition and 3.4 s for the 6 s initial TTA condition) before decelerating more to come to a stop 2 to 2.5 meters from the participant. These brake onset timings were designed based on the previous experiment's mixed manoeuvres, ensuring 1) that the deceleration rates of the 2-stage deceleration scenarios were greater than the deceleration rates of the corresponding 1-stage deceleration scenarios and 2) the $\dot{\tau}$-level similarity of corresponding segments. The brake onset timings design was finalised by qualitatively matching the previous with the patterns observed in the speed profiles from the informal static driving simulator tests.

- Slowing Early and Approaching: Early Deceleration and Continuation without Stopping Manoeuvre

This manoeuvre involves an early reduction in speed, maintained as the vehicle approaches the pedestrian, without a complete stop. This behaviour can signal to pedestrians that they are acknowledged and can safely cross (Risto et al., 2017). The strength of this signal, however, is likely linked to the deceleration rate and the distance from the crossing (Domeyer et al., 2019). This short stopping or early slowing strategy represents a subtle form of yielding, potentially balancing driver convenience with pedestrian safety. This manoeuvre design was based on the qualitative speed profiles obtained by the informal static driving simulator tests, since it was not tested in the previous experiment and there were no data available in the literature and will be referred to as Short Slowing. In the short slowing manoeuvre scenarios, the vehicle combined an initial phase of sharper deceleration and then a second phase where it maintained constant speed (lower than the initial speed), to, again, explore the impact of the vehicle's behaviour change, during the simulation, on the pedestrian's belief updating Initially, it was decided to use the same timings as the brake onsets of the Two-stage deceleration manoeuvres. However, an informal static driving simulator test and early experimental design tests in HIKER revealed that these timings were quite short in this case, leading to high and unrealistic deceleration rates. For that reason, it was deemed necessary to prolong the decelerating phases' durations,

leading to 1.8 s for the 3 s initial TTA condition and 3.6 s for the 6 s initial TTA condition, which could be still described as exaggerated but within realistic levels.

It is worth noting that that not all 3 different initial speeds were used for all the vehicle manoeuvres. All initial speeds (20, 40 and 60 km/h) were used for the *'validation'* manoeuvres (Constant speed and Deceleration), while only the extrapolated initial speeds (20 and 60 km/h) were used for the *'generalising'* manoeuvres (Two-stage deceleration and Short slowing). The decision to not include the 40 km/h initial speed in the 'Two-stage Deceleration' and 'Short Slowing' manoeuvres was made to manage the overall duration of the experiment (including all possible combinations of the independent variables would have led to extremely long experimental sessions for the participants).

In the designed scenarios, the deceleration rates were between 0.12 ~ 4.83 m/s^2. The lower deceleration rates that were used in the experiment are frequently observed when vehicles normally yield to pedestrians (Carlowitz et al., 2024; Yang et al., 2024), while some of the greater deceleration rates are observed when vehicles slow down before intersections on roads with higher speed limits or in emergency situations, for example a pedestrian abruptly stepping on the road (Bokare and Maurya, 2017; Kudarauskas, 2007). The deceleration rates range of the 'Deceleration' scenarios was wider in this experiment than the respective range in the previous experiment. Similarly, the deceleration rates range of the 'Two-stage Deceleration' scenarios was also wider in this experiment than the respective range of the 'Mixed' scenarios of the previous experiment. Including this wide range of deceleration rates in the experimental design, allowed testing the models' ability to predict pedestrian beliefs under both common and more extreme, potentially more ambiguous, conditions, when driver behaviours are less typical and/or more urgent. Table 4.1 shows the kinematic parameters of the designed traffic scenarios of the approaching vehicle. In summary, the experiment consisted of 4 types of manoeuvres, 3 initial speeds, and 2 initial TTAs, resulting in a total of 20 conditions.

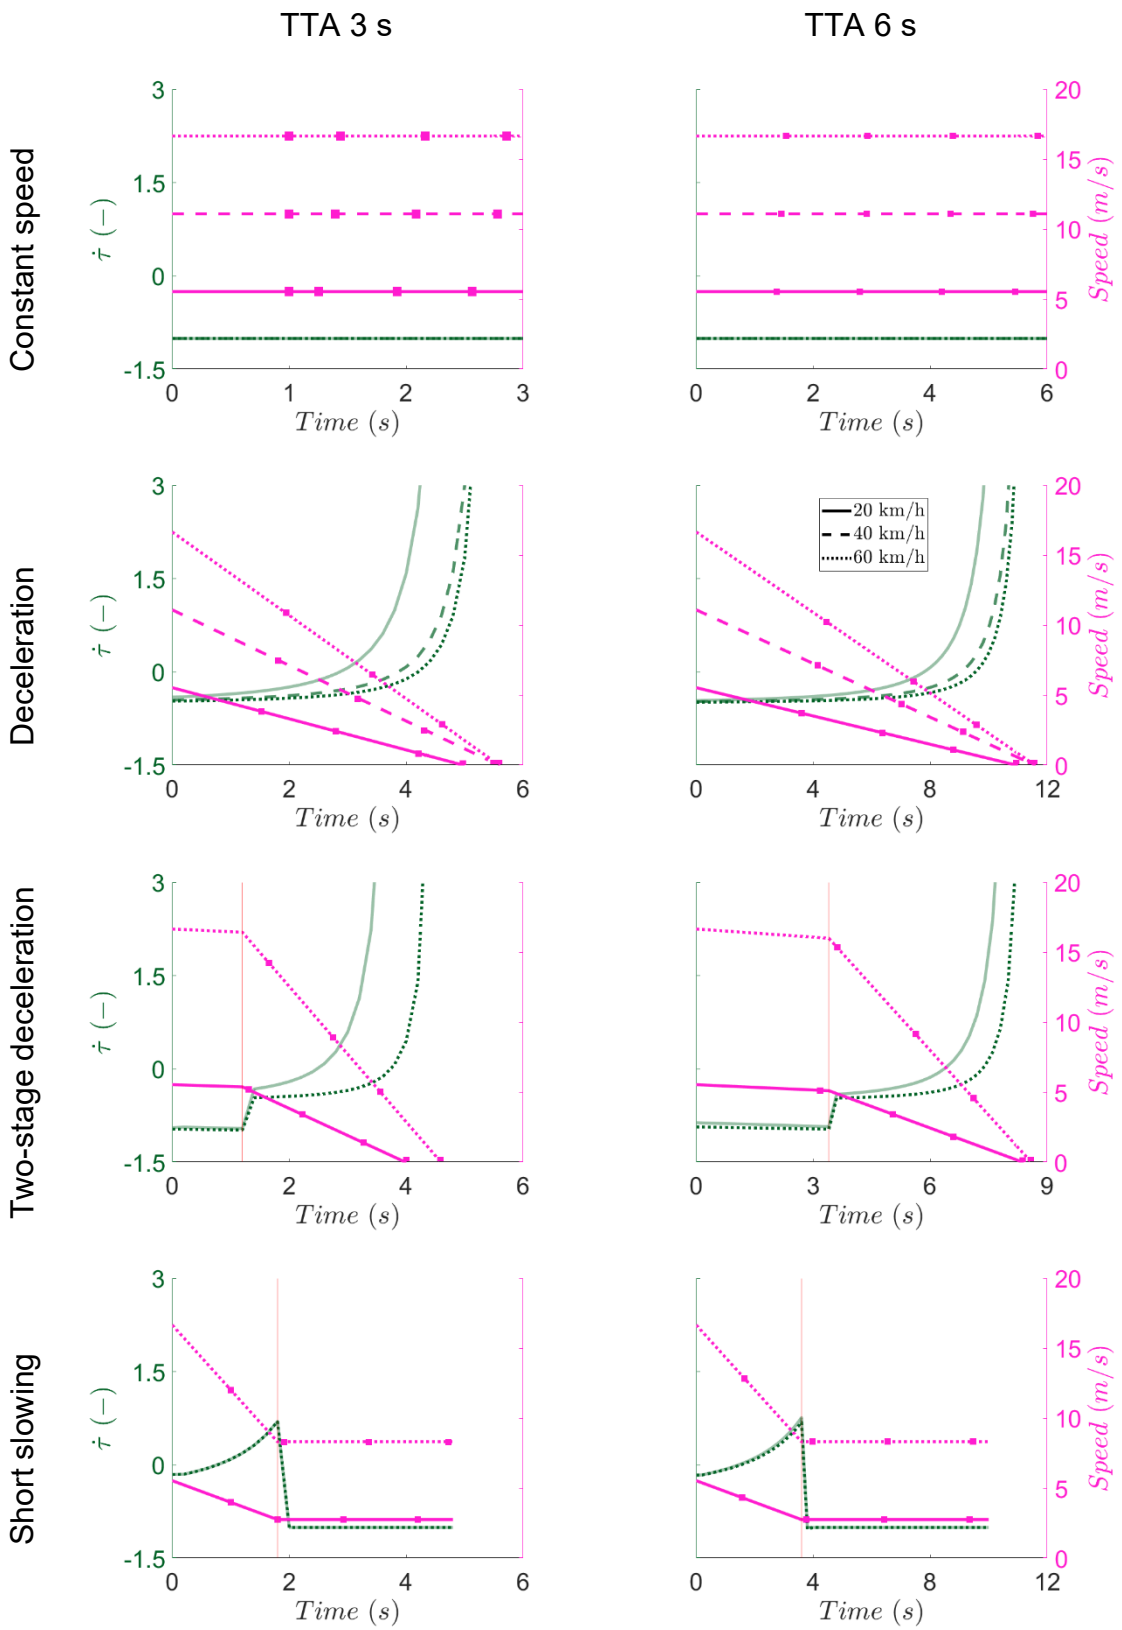

Figure 4.1 - Vehicle kinematics of all 20 scenarios. The vehicle's speed profile is denoted using the pink curves and the respective $\dot{\tau}$ time history by the dark green. The red vertical lines in the Two-stage deceleration and Short slowing manoeuvres indicate the timings of the harsher brake onset and the brake release respectively

Table 4.1 - Parameters for traffic scenarios

Manoeuvre	Initial TTA (s)	Initial speed (km/h)	Initial distance to pedestrian (m)	Deceleration rates (m/s^2)	Constant speed phase duration (s)
Constant speed	3	20	16.67	-	[0, 3]
		40	33.33	-	[0, 3]
		60	50	-	[0, 3]
	6	20	33.33	-	[0, 6]
		40	66.67	-	[0, 6]
		60	100	-	[0, 6]
Deceleration	3	20	16.67	1.11	-
		40	33.33	1.98	-
		60	50	2.98	-
	6	20	33.33	0.51	-
		40	66.67	0.97	-
		60	100	1.44	-
Two-stage deceleration	3	20	16.67	0.13 & 1.93	-
		60	50	0.20 & 4.83	-
	6	20	33.33	0.12 & 1.03	-
		60	100	0.19 & 3.08	-
Short slowing	3	20	16.67	1.54	(1.8, 4.8]
		60	50	4.63	(1.8, 4.8]
	6	20	33.33	0.77	(3.6, 10]
		60	100	2.31	(3.6, 10]

4.1.4 Tasks and procedure

In essence, the core tasks and procedure were maintained from Chapter 2, with minor adjustments to the order of the tasks and the segments calculations in order to accommodate the kinematics of the new experimental design and the expanded scenario set and maintain participant engagement. In the previous experiment, the participants first completed the two road crossing blocks and then they completed the two behaviour estimation blocks. In the new experiment, even though there were again four main experimental blocks (two road crossing and two behaviour estimation), they were alternated according to the following sequence: Road crossing (Block 1) → Behaviour estimation (Block 2) → Road crossing (Block 3) → Behaviour estimation (Block 4), as seen in Figure 4.2, to better balance participant fatigue throughout the experiment, given the increased number of scenarios and segments in this study, as the road crossing blocks included a more active task in comparison to the monotonous behaviour estimation task.

The road crossing task, consistent with Chapter 2, required participants to cross the road if and when they felt safe to do so, while interacting with an approaching (virtual) vehicle. The instructions given to the participants were the same as the ones described in the previous experiment. Block 1 was preceded by a practice session with 10 trials for the participants to familiarise themselves with the task and the environment. The road crossing blocks (1 and 3) included the same 20 experimental conditions, that were randomised within each of these two blocks. The initiation and procedure of the trials of the road crossing task were exactly the same to the previous experiment, as described in detail in Subsection 2.1.4. After completing each of the road crossing blocks, a short break was taken, before the respective following behaviour estimation block began.

The behaviour estimation task, also consistent with Chapter 2, required participants to judge whether the vehicle was stopping or not stopping. As in Chapter 2, each of the scenarios was truncated into four segments of varying lengths, following the paradigm of Baker et al. (2009). The segmentation aimed to provide increasing visual cues and overall behaviour evidence, in sequence, during a scenario.

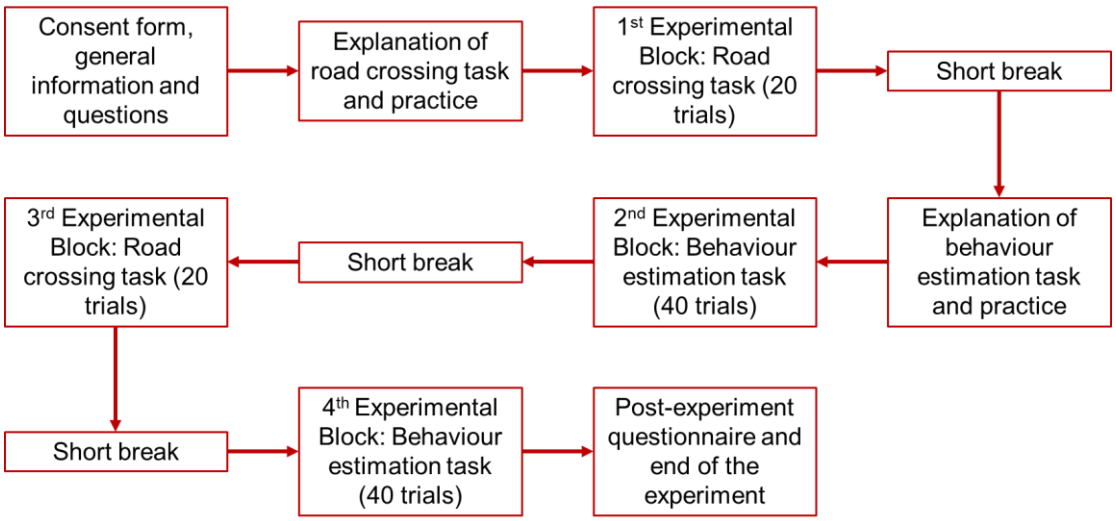

Figure 4.2 - Procedure of the experiment

However, the calculation of these segment divisions differed slightly from Chapter 2 due to the overall experimental design changes and more so due to the inclusion of the 'Short Slowing' manoeuvre. Specifically, the 'Constant Speed' and 'Short Slowing' scenarios were divided using the linear division method, as the visual cue for collision judgment remains mostly constant in these scenarios, as was illustrated in Figure 4.1. The 'Deceleration' and 'Two-stage Deceleration' scenarios were divided using the logarithmic division method. The exact equations for calculating the division were similar to the equations in Subsection 2.1.4, with the linear division equation now applied also to the short slowing scenarios. Similarly to the previous experiment, there were cases where the duration of the 1st segment as calculated by this method was too short, so in order to avoid these very subtle stimuli, their duration was fixed at 1 s. According to the above division methods, the traffic scenarios of 20 experimental conditions were divided into 80 segments in total (indicated as pink square markers on the speed profiles in Figure 4.1).

The initiation procedure of the behaviour estimation trials replicated the initiation procedure of the respective task of the previous experiment. After each segment presentation, the environment was obscured, and the same questions that were described in Chapter 2, regarding the vehicle's behaviour and the participant's confidence, followed. So, the participants only had to observe the traffic scenarios and answer the questions, while no road crossing was required. The first behaviour estimation block (Block 2 in the experiment; see Figure 4.2) was preceded by a practice session with 10 trials. The formal behaviour estimation blocks had a total of 80 segments (40 in Block 2 and 40 in Block 4),

which were presented in an order that was randomised per participant, meaning that each participant experienced each segment of each scenario once.

4.2 Model definitions

As a reminder, the models being evaluated in this chapter are the same as in the previous chapter: the observation-based model (Ob), the value-based model (Vb), and the combined observation and value-based model (Ob+Vb). Ob predicts pedestrian beliefs solely based on observed vehicle behaviour, more specifically $\dot{\tau}$ values, (same model that was presented in Chapter 2 and 3). Vb, which was introduced in Chapter 3, on the other hand, predicts pedestrian beliefs based on the assumption that pedestrians infer vehicle intentions by considering what actions would be most rational or value-maximising for the driver. Finally, Ob+Vb integrates both observation-based cues and value-based reasoning.

It is important to clarify that while the new experiment included four distinct vehicle manoeuvres, the models aimed to capture the pedestrian's subjective beliefs regarding two mutually exclusive perceived vehicle behaviours, i.e., the vehicle decelerating with the intention to stop for them (P_s) or the vehicle maintaining constant speed with no intention to stop for them (P_{ns}).

4.3 Model fitting

This section describes the model fitting approaches of the previously developed behaviour estimation models to the new dataset acquired collected as described above. The model fitting followed the same methodology as described in Chapter 3, employing a global derivative-free optimisation algorithm (Section 3.2) to optimise model parameters. The following three model fitting approaches were used:

First, to provide a baseline for comparison and assess the generalisability of the models, model predictions generated using the optimal parameter combination identified in Chapter 3 were compared against the Chapter 4 data.

This allowed for an initial evaluation of how well the models, trained on the previous dataset, generalised to the new experimental scenarios.

Second, a focused model fitting exercise was performed, specifically adjusting the added margin parameter denoted as M. This parameter represents the strength of the prior belief towards decelerating behaviour, influencing the value-based component (Vb) and, consequently, the augmented model (Ob+Vb). The rationale behind this targeted adjustment stemmed from the observation that the new experimental scenarios in Chapter 4 might elicit different prior beliefs in participants compared to those predicted in Chapter 3. For instance, the inclusion of 'Short Slowing' manoeuvres and the differences between 'Two-stage Deceleration' and 'Mixed' manoeuvres could alter participants' expectations of vehicle behaviour. Therefore, the value of M was adjusted from the optimal value found in Chapter 3 to a new value, while keeping the rest of the parameter settings constant. This allowed us to explore the sensitivity of the models to this parameter within the context of the new dataset, and to determine whether the prior belief strength needed to be recalibrated to better reflect the participants' inferences in these new experimental conditions.

Finally, a complete model fitting procedure was undertaken. All model parameters were optimised on the Chapter 4 dataset, again using the methodology described in Chapter 3. This comprehensive fitting process aimed to maximise the predictive accuracy of each model on the new data, providing a benchmark of their performance under optimal conditions for the current experimental setup. The results of these fitting procedures, including the optimised parameter values and the resulting model fit metrics, are presented and discussed in the subsequent sections of this chapter.

4.4 Model selection

The model selection procedure was the same as the one described in Chapter 3. A quantitative approach was employed to assess the predictive accuracy of the observation-based, value-based, and augmented models on the Chapter 4 dataset. Specifically, Root Mean Squared Errors (RMSEs) were calculated between model predictions and participant ratings. To compare the predictive

accuracy of the three models, Bootstrap Cross-Validation (BSCV; Baker et al., 2009; Cohen et al., 1955), a robust, non-parametric technique for model selection, was utilised.

4.5 Results

This section presents empirical observations, accompanied with statistical analysis, from the experiment and then several analyses testing the predictive accuracy of the proposed models – Observation-based (Ob), Value-based (Vb), and the combined Observation and Value-based model (Ob+Vb) – in capturing pedestrians' inferences about an approaching vehicle's behaviour, based on the data collected in the experiment described in Section 4.1.

4.5.1 Pedestrian beliefs regarding the vehicle's behaviour

In this subsection the overall trends of the of the participant's judgements are presented. As described in Sections 2.1 and 4.1, the participants gave two answers during the behaviour estimation tasks of the experiments, regarding which behaviour they believed the vehicle was exhibiting (binary choice between stopping and maintaining speed) and how certain they were about that (subjective rating on a 1-9 Likert scale). These two answers were transformed linearly into a probability P_s denoting the pedestrian belief that the approaching vehicle was exhibiting the stopping (decelerating to stop) behaviour. Notably, $P_s = 0$ would mean that the pedestrian fully believes that the vehicle was exhibiting the other possible behaviour, maintaining constant speed with no intention to stop, $P_{ns} = 1$, since the two behaviours are mutually exclusive and exhaustive (in both presented experiments), $P_s + P_{ns} = 1$.

The pedestrian beliefs regarding the approaching vehicle's behaviour were analysed using a mixed-effects linear regression model, instead of a 4-way ANOVA, as presented in Chapter 2, due to the unbalanced design of the new experiment (i.e., not all initial speeds were tested for all four manoeuvres). The pedestrians' belief that the vehicle exhibited the stopping behaviour was the dependent variable. The segment, the vehicle's initial speed, the initial TTA and the driving manoeuvre were considered to be the fixed effects. Lastly, the participants' individual differences were modelled as a random intercept. A model

incorporating both a random intercept and random slopes (for all main fixed effects) was also evaluated. While the maximal random effects model yielded a lower Akaike Information Criterion (AIC) value compared to the random intercept-only model, the Bayesian Information Criterion (BIC) value it yielded was higher. The latter indicates that the substantial increase in model complexity from adding random slopes for all main predictors, on top of the random intercept, was not justified by a sufficient improvement in fit for this dataset. Therefore, the more parsimonious random intercept-only model, which effectively accounts for baseline inter-participant variability, was used for the final analysis.

The analysed mixed-effects linear regression model is presented below, using the Wilkinson notation (Wilkinson and Rogers, 1973). Equation (4.1) indicates that the Belief could be predicted by Segment, Manoeuvre, Speed and TTA, along with their interactions. Then, the last term represents the random intercept accounting for the inter-individual variability in participants' (PID) average level of Belief, meaning that they might be more cautious or more trusting than other individuals when judging if an approaching vehicle is stopping or not.

$$Belief \sim Segment * Manoeuvre * Speed * TTA + (1|PID) \quad (4.1)$$

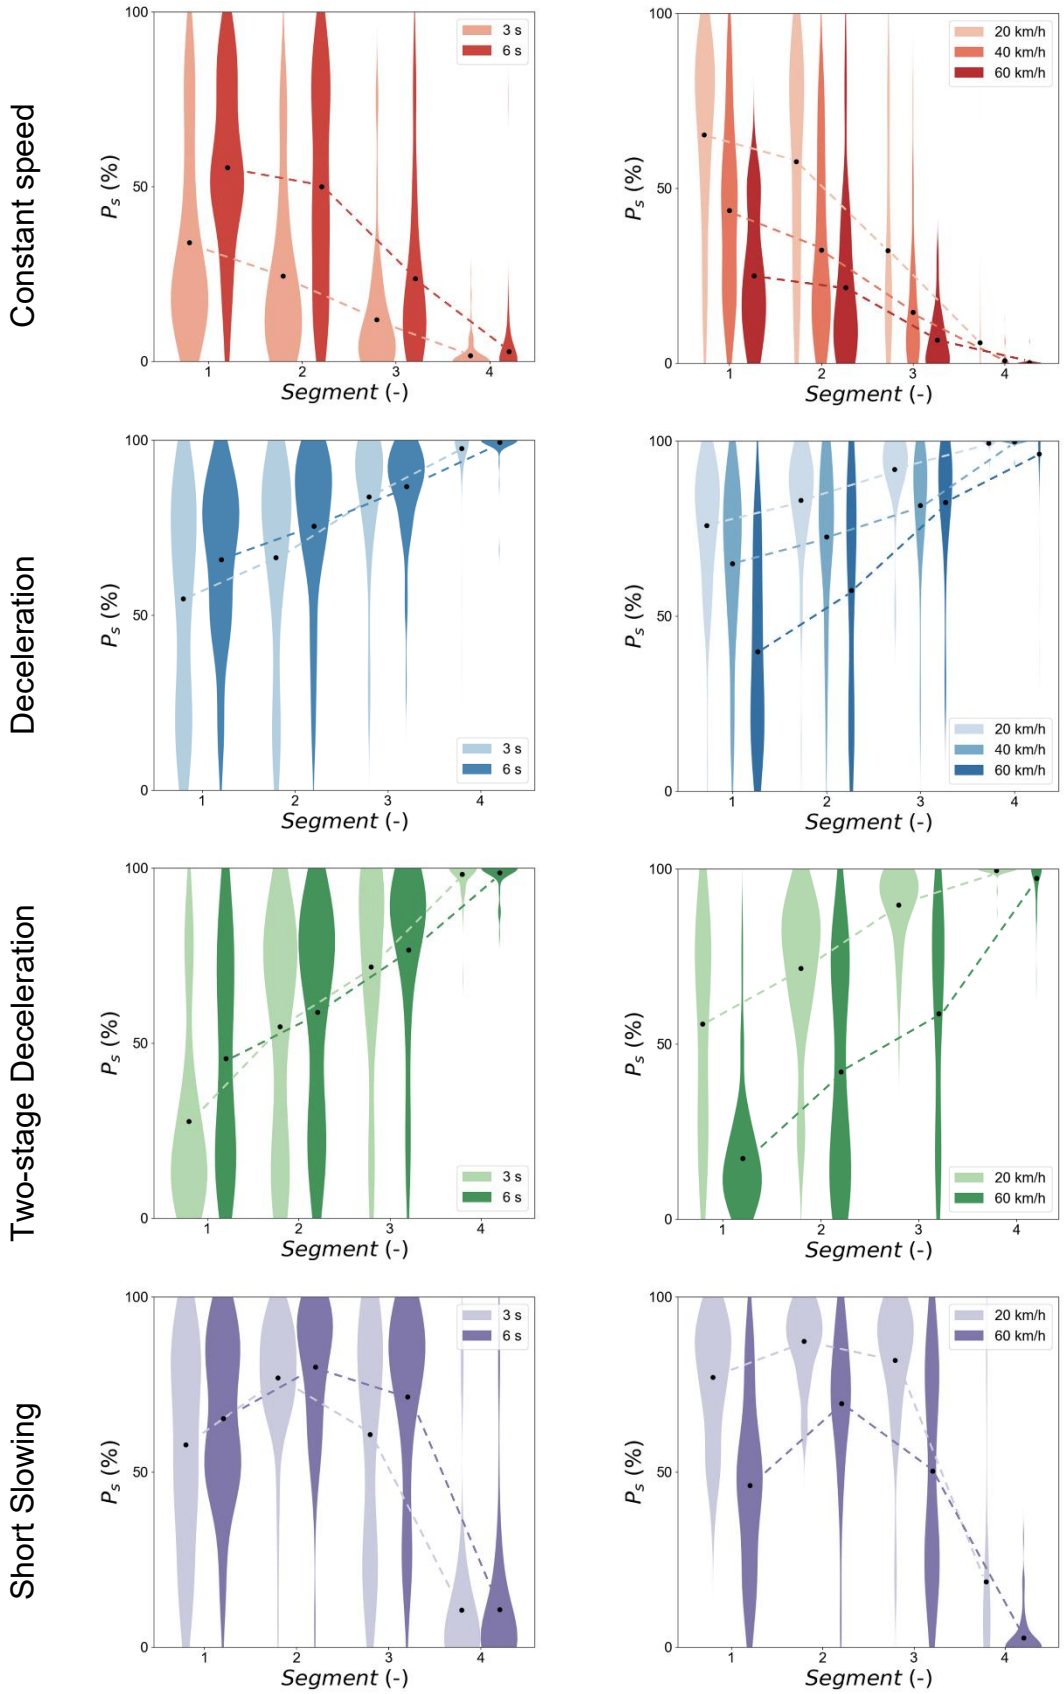

Figure 4.3 - Pedestrians' beliefs regarding the approaching vehicle's behaviour as a function of segment, initial TTA (left column) and initial speed of the vehicle (right column)

The full results of the mixed-effects linear regression model can be found in Appendix A . The fixed effects coefficients reveal several statistically significant predictors of pedestrian beliefs. Notably, 'Speed' exhibited a significant negative effect (Estimate = -1.169, $p < 0.001$), indicating that as the initial speed of the vehicle increased the likelihood of pedestrians believing the vehicle would yield decreased. This effect can be clearly seen in the panels of the right column of Figure 4.3. Conversely, 'TTA' showed a significant positive effect (Estimate = 13.137, $p < 0.001$), suggesting that longer initial TTA was associated with a higher probability of pedestrians believing the vehicle would stop. This effect can be clearly seen in the panels of the left column of Figure 4.3. The two findings described above were consistent with the results of the 4-way ANOVA (Subsection 2.5.1.2), which was performed on the data collected from the first experiment.

The vehicle's ‘Manoeuvre’ also played a significant role in shaping pedestrian beliefs. The 'Short slowing' manoeuvre was selected as the reference category for these comparisons, given that it was the newest and most distinct driving manoeuvre relative to those tested in the previous experiment. ‘Constant speed’ manoeuvres had a significant negative effect (Estimate = -52.852, $p < 0.001$), indicating that pedestrians were significantly less likely to believe the vehicle would stop during a ‘Constant speed’ manoeuvre compared to a vehicle performing a ‘Short slowing’ manoeuvre (reference category). This could be because in the beginning of the ‘Short slowing’ manoeuvres the vehicle was actually stopping, thus shifting the pedestrians’ beliefs more towards the stopping behaviour than in the beginning of ‘Constant speed’ manoeuvres – for comparison see the red and purple trends in Figure 4.4. Conversely, ‘Deceleration’ manoeuvres had a significant positive effect (Estimate = 34.857, $p < 0.001$), suggesting that pedestrians were significantly more likely to believe the vehicle would stop during a ‘Deceleration manoeuvre’ compared to a vehicle performing a ‘Short slowing’ manoeuvre (blue and purple trends in Figure 4.4). Similarly to ‘Deceleration manoeuvres’, in ‘Two-stage deceleration’ manoeuvres, pedestrians tended to believe that the vehicle was stopping more, in comparison to their beliefs during a ‘Short slowing’ manoeuvre (green and purple trends in Figure 4.4), supported by a significant positive effect (Estimate = 11.784, $p < 0.05$). Finally, the model was rerun using the ‘Two-stage deceleration’ manoeuvre

as the reference category allowing the comparison between the ‘Two-stage deceleration’ and ‘Deceleration’ manoeuvre categories. ‘Deceleration’ manoeuvres had a significant positive effect (Estimate = 73.299, $p < 0.001$), suggesting that pedestrians’ beliefs were significantly more likely to lean towards the vehicle stopping during ‘Deceleration’ manoeuvres compared to a vehicle performing a ‘Two-stage deceleration’ manoeuvre (blue and green trends in Figure 4.4).

Several significant interaction effects were also observed, indicating that the influence of one predictor on pedestrian belief depended on the level of another. For instance, the significant negative interaction for Segment 1 and ‘Speed’ (Estimate = -0.283, $p < 0.05$) suggests that the negative impact of speed on believing that the vehicle was yielding was stronger in Segment 1 compared to the reference segment (Segment 4). The significant positive interaction between Speed and TTA (Estimate = 0.027, $p < 0.05$) suggests a complex relationship where the effect of TTA on belief might be moderated by the vehicle's speed (Figure 4.3). Furthermore, the significant interaction effects underscore that these main effects might not be independent but rather contingent on the specific context defined by the vehicle's segment of approach, speed, and time-to-arrival. Finally, the random intercept for the inter-participant differences had a standard deviation of 3.896, indicating significant variability in baseline beliefs across individuals.

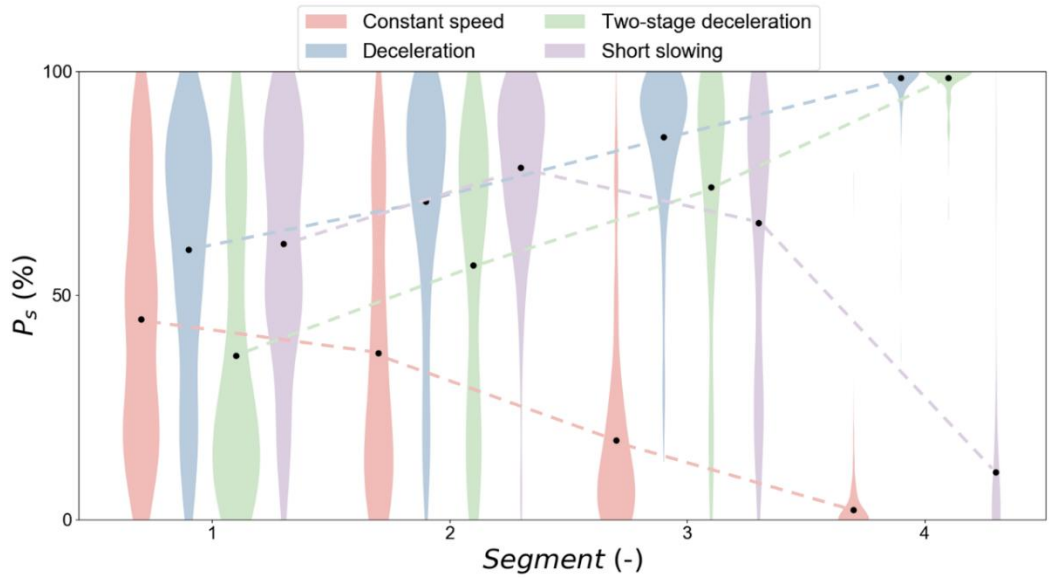

Figure 4.4 - Pedestrians’ beliefs regarding the approaching vehicle’s behaviour as a function of judgement point and the vehicle’s driving manoeuvre. Means are represented by black dots and dashed lines indicate the trends of the average beliefs

Due to the differences in the pedestrians' beliefs across the different manoeuvres (Figure 4.4), as explained above, it was deemed necessary to divide the data per driving manoeuvre and analyse them separately with the following mixed-effects linear regression model (Equation (4.2)).

$$Belief \sim Segment * Speed * TTA + (1|PID) \quad (4.2)$$

In the 'Constant speed' manoeuvres, pedestrians were more likely to believe that the vehicle was stopping when the vehicle had a lower initial speed and the initial TTA was larger (Estimate = -1.169, $p < 0.001$ and Estimate = 13.137, $p < 0.001$) – both effects are apparent in the "red" row panels of Figure 4.3. Additionally, as time was passing and the vehicle got closer to them, pedestrians were less likely to believe that the vehicle was stopping, indicated by the segments having a significant negative effect (Estimate = -15.576, $p < 0.01$).

In the 'Deceleration' manoeuvres, pedestrians were more likely to believe that the vehicle was stopping only when the vehicle had a lower initial speed (Estimate = -2.072, $p < 0.001$). On the other hand, the main effects of initial TTA and segments were not statistically significant, and especially the latter is clear in the left "blue" panel of Figure 4.3.

In the 'Two-stage deceleration' manoeuvres, pedestrians were more likely to believe that the vehicle was stopping when the vehicle had a lower initial speed (Estimate = -0.932, $p = 0.01$). As more time passed, and the vehicle got closer, pedestrians were more likely to believe that the vehicle was stopping (Estimate = 24.292, $p < 0.001$). On the other hand, the main effect of initial TTA was not statistically significant and can be seen in the left "green" panel of Figure 4.3.

In the 'Short slowing' manoeuvres, pedestrians were more likely to believe that the vehicle was stopping when the vehicle had a lower initial speed (Estimate = -1.081, $p < 0.05$). As more time passed, and the vehicle got closer, pedestrians were more likely to believe that the vehicle was not stopping (Estimate = -18.708, $p < 0.05$). On the other hand, the main effect of initial TTA was not statistically significant and can be seen in the left "purple" panel of Figure 4.3.

Besides the specific statistical effects within each manoeuvre type, visual inspection of the average pedestrian belief trends suggested further nuanced patterns. The average beliefs of the ‘Two-stage manoeuvres’ seemed to be a combination of the average beliefs from the first segment of the ‘Constant speed’ manoeuvres and of the later segments of the ‘Deceleration’ manoeuvres. In a similar manner, the average beliefs of the ‘Short slowing’ manoeuvres seemed to be a combination of the average beliefs of the first segment of the ‘Deceleration’ manoeuvres and of the later segments of the ‘Constant speed’ manoeuvres. To assess these qualitative observations, separate mixed-effects linear regression models were applied. These separate analyses supported the above observations by not finding any significant differences when comparing the pedestrian beliefs of the first and later segments between 1) Constant speed and Two-stage deceleration manoeuvres and 2) Deceleration and Short slowing.

4.5.2 Model predictions vs subjective ratings

The analysis investigated the predictive accuracy of the three distinct models in estimating pedestrian's average subjective beliefs. This was accomplished by assessing both correlational (Spearman's rank correlation) and error-based (RMSE) metrics. The predictive outputs of the models presented in this subsection were derived using the three model fitting approaches described in Section 4.3:

1) Using the optimal parameter settings found in the analysis of Chapter 3.
2) Using the optimal parameter settings found in the analysis of Chapter 3 but fitting only the added margin M parameter to the new dataset.
3) Fitting all model parameters to the new dataset.

Ob	Vb	Ob+Vb
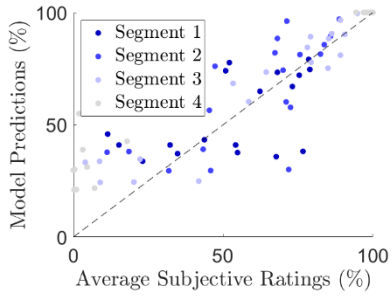	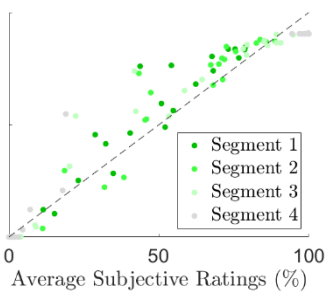	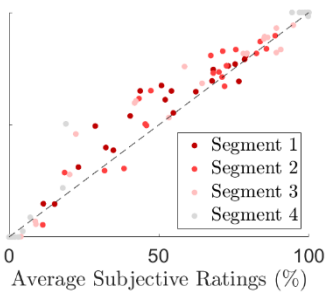
$\rho_S = 0.86$	$\rho_S = 0.98$	$\rho_S = 0.98$
$RMSE = 18.99$	$RMSE = 10.61$	$RMSE = 8.73$
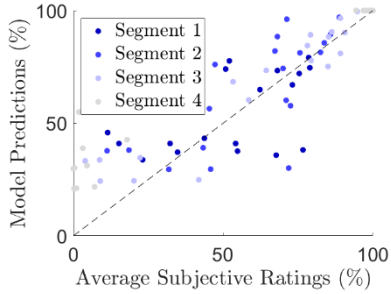	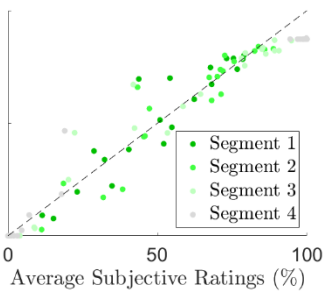	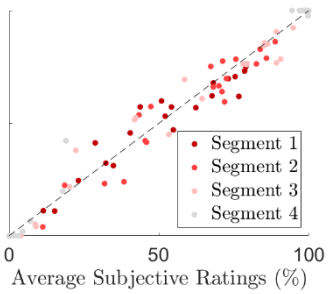
$\rho_S = 0.86$	$\rho_S = 0.98$	$\rho_S = 0.98$
$RMSE = 18.99$	$RMSE = 9.19$	$RMSE = 6.67$
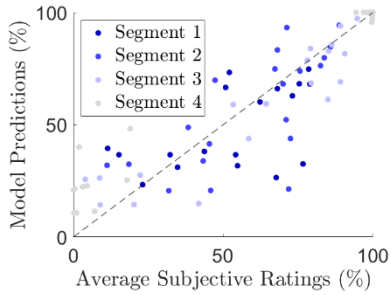	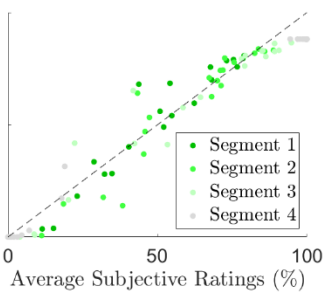	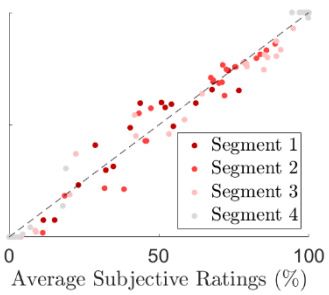
$\rho_S = 0.86$	$\rho_S = 0.98$	$\rho_S = 0.98$
$RMSE = 16.84$	$RMSE = 8.41$	$RMSE = 6.00$

Figure 4.5 - Scatter plots of model predictions using best-fitting parameter settings (y-axes) versus pedestrians' average beliefs (x-axes) for all 20 kinematic scenarios of the approaching vehicle. The three columns indicate the tested model (Ob, Vb and Ob+Vb). The three rows indicate the model fitting approach that was used (can be found in the enumeration of the previous page)

Figure 4.5 presents scatter plots comparing participant judgments with predictions from the Ob, Vb, and Ob+Vb models – which are represented by blue, green and red colours respectively, similarly to the previous chapter, illustrating their respective accuracies. In the first-row panels, the models were parameterised according to the best-fitting values found in the analysis of Chapter

3 (as detailed in Table 3.2). Consistent with the findings from Chapter 3, visual inspection of Figure 4.5 reveals that Vb and Ob+Vb again exhibited fewer poorly predicted datapoints compared to Ob. This observation is further supported by the obtained correlation and RMSE values. Notably, Vb and Ob+Vb demonstrate the same near-perfect positive monotonic rank association (Spearman's rank correlation ≈ 1) between pedestrian average beliefs and model predictions as was observed in the analysis of Chapter 3. In contrast, Ob shows a lower Spearman's rank correlation on this new dataset than Vb and Ob+Vb, a pattern which was observed also in the previous chapter. Furthermore, the scatter points for Vb and Ob+Vb cluster closely along the identity line (Pearson's rho ≈ 1), again mirroring the near-perfect positive linear relationship found previously and suggesting that their predictions closely approximate mean pedestrian beliefs. Finally, considering the Root Mean Squared Error (RMSE), a measure of proximity to the identity line where lower values indicate better predictive performance, Ob+Vb demonstrates a lower RMSE than Vb, while Ob exhibits the highest error, consistent with the findings of Chapter 3.

Importantly, while the Spearman's rank correlations remained identical to those found in Chapter 3, except for Ob's (Ob: 0.89→0.86, Vb: 0.98→0.98 and Ob+Vb: 0.98→0.98), the RMSE values for all three models are slightly higher on this new dataset. The decrease of Ob's rank correlation could be due to the addition of more segments during which the approaching vehicle maintained constant speed (i.e., the later segments of 'Short slowing' manoeuvres), and since Ob's predictions were poor in such cases. The RMSE increase for all three models could be due to testing the models on this new dataset, which was an extrapolated version of the previous dataset, in terms of initial speeds and driving manoeuvres. Partial and/or full refitting slightly improved the predictive performance of the models, bringing their RMSE values very close to the RMSE values obtained in Chapter 3.

Considering that the models were applied to a completely new set of pedestrian participants experiencing the new experimental design (with new untested kinematic scenarios), the level of performance achieved, particularly by Ob+Vb, could be argued to show a reasonably good degree of generalisation. The models were able to maintain a strong relationship with the pedestrian's subjective beliefs even when presented with novel data.

Following this general analysis of the models' predictive performance, more detailed analyses of the models' behaviour are presented in Subsections 4.5.2.1, 4.5.2.2 and 4.5.2.3. In each one of these subsections, the parameter settings of the models were obtained by the three model fitting approaches described in Section 4.3.

4.5.2.1 Parameter settings of Chapter 3

An analysis comparing pedestrian data and model predictions across the new kinematic scenarios elucidates the influence of varying vehicle kinematic conditions on pedestrian belief updating regarding the approaching vehicle's behaviour and assesses the fidelity with which the proposed models capture the average pedestrian beliefs.

Specifically, vehicle manoeuvres are categorised and presented in the following figures: Figure 4.6, Figure 4.7, Figure 4.8 , and Figure 4.9. Figures 4.5 and 4.6 are including the scenarios with 40 km/h initial vehicle speed for both initial TTAs (3 and 6 s). The average pedestrian beliefs of the last experiment are shown in these four specific scenarios. These exact kinematic scenarios were tested in the previous experiment as well, so they could serve as scenarios through which model and experiment validation were tested. The rest of the kinematic scenarios were illustrated in the abovementioned figures, aiming to present all the new scenarios of the new experiment, so they could serve as scenarios through which model generalisability was tested.

These 20 plots (one for each kinematic scenario) illustrate general patterns in pedestrian beliefs as predicted by the three models. Within Figure 4.6, Figure 4.7, Figure 4.8 and Figure 4.9 (and all the rest figures of this chapter), the predictions of the Ob, Vb, and Ob+Vb models are represented by blue, green, and red curves, respectively, as was also the case in Chapter 3. The average pedestrian beliefs are depicted as standard error of the mean bars of the pedestrians' beliefs (the black ones indicate the average beliefs from the new dataset, while the grey ones indicate the average beliefs from the previous dataset) at the pre-defined judgment timings (as detailed by the pink square marks in Figure 4.1). For clarification, the model prediction curves in Figure 4.6, Figure 4.7, Figure 4.8 and Figure 4.9 are exactly the same as the ones produced in Chapter 3. With a quick visual inspection of the 40 km/h scenarios in Figure 4.6 and Figure 4.7 (second row panels) it can be seen that the average pedestrian beliefs of the new

experiment are consistent and quite close to the ones of the previous experiment, with a slight tendency of the average beliefs of the new experiment to be leaning a bit more towards the non-stopping behaviour (lower P_s or equally higher P_{ns}).

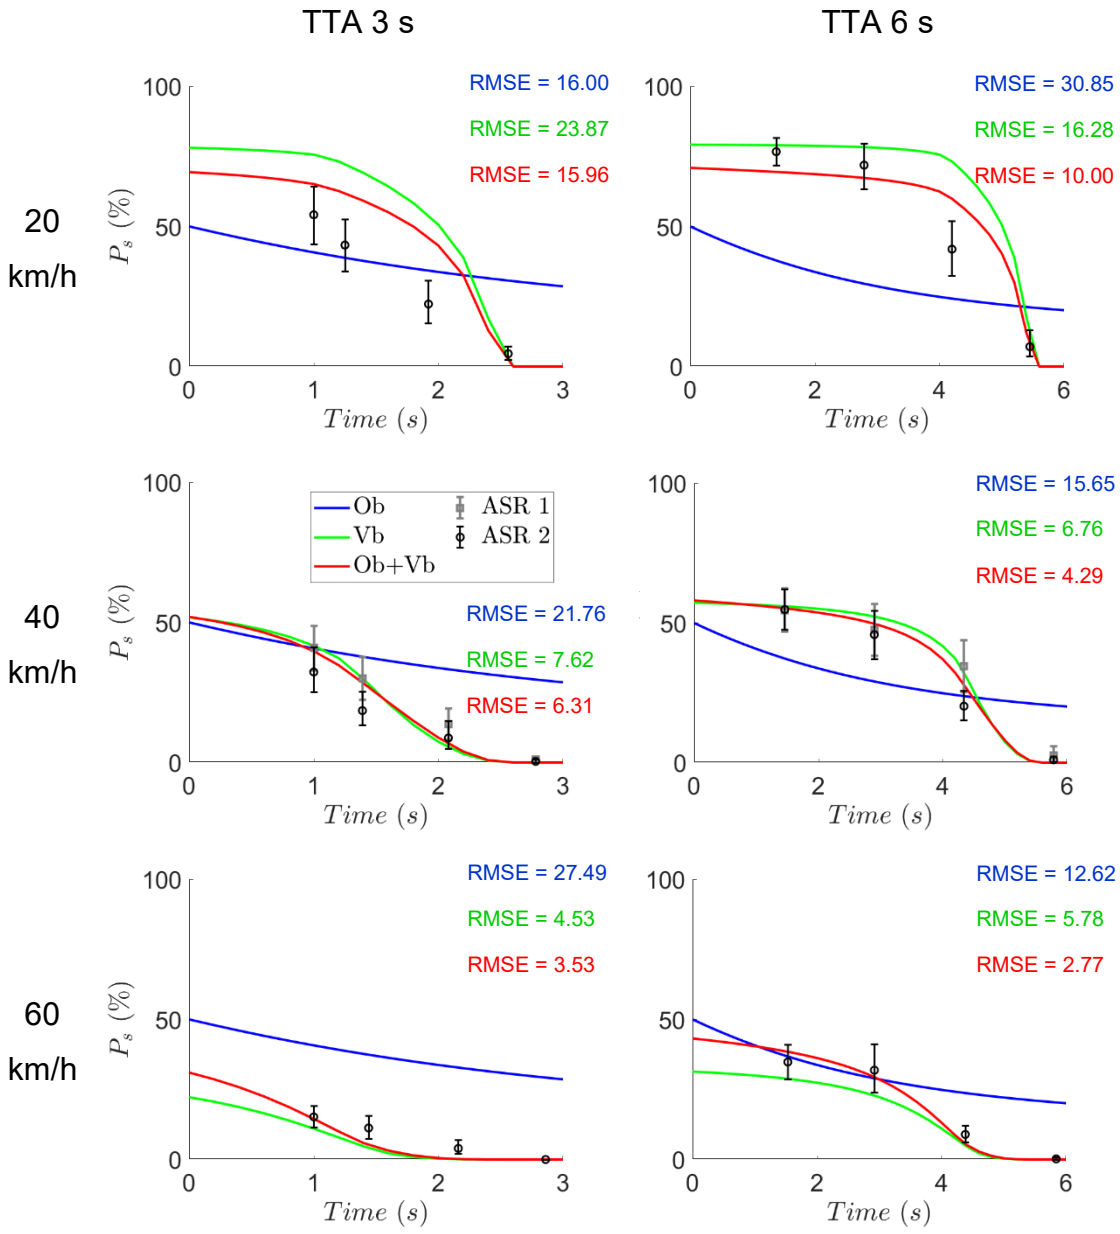

Figure 4.6 - Comparison between model predictions and average pedestrian beliefs for the constant speed manoeuvres. The RMSEs are regarding the new dataset only (i.e., Average Subjective Ratings 2)

In the constant speed scenarios (Figure 4.6), similarly to Chapter 3, Vb and Ob+Vb outperformed Ob in capturing the pedestrian belief patterns. Ob consistently predicted initial 50/50 belief uncertainty, while Vb and Ob+Vb reflected the influence of speed and TTA. Ob's reliance on $\dot{\tau}$ limited its ability to differentiate beliefs across scenarios, a limitation addressed by Vb's and Ob+Vb's incorporation of additional information. Higher vehicle speeds correlated with

stronger predicted beliefs in non-stopping behaviour, while larger TTAs correlated with weaker such beliefs. These patterns can be further explained by examining the models' evidence sources (Subsections 3.4.2 and 4.5.3). Ob+Vb's predictions in the 40km/h scenarios support the successful validation of the suggested behaviour estimation mechanisms and the experimental paradigm. In addition, Vb and Ob+Vb seemed to be able to generalise quite well (especially for the higher speeds) to the new initial vehicle speeds.

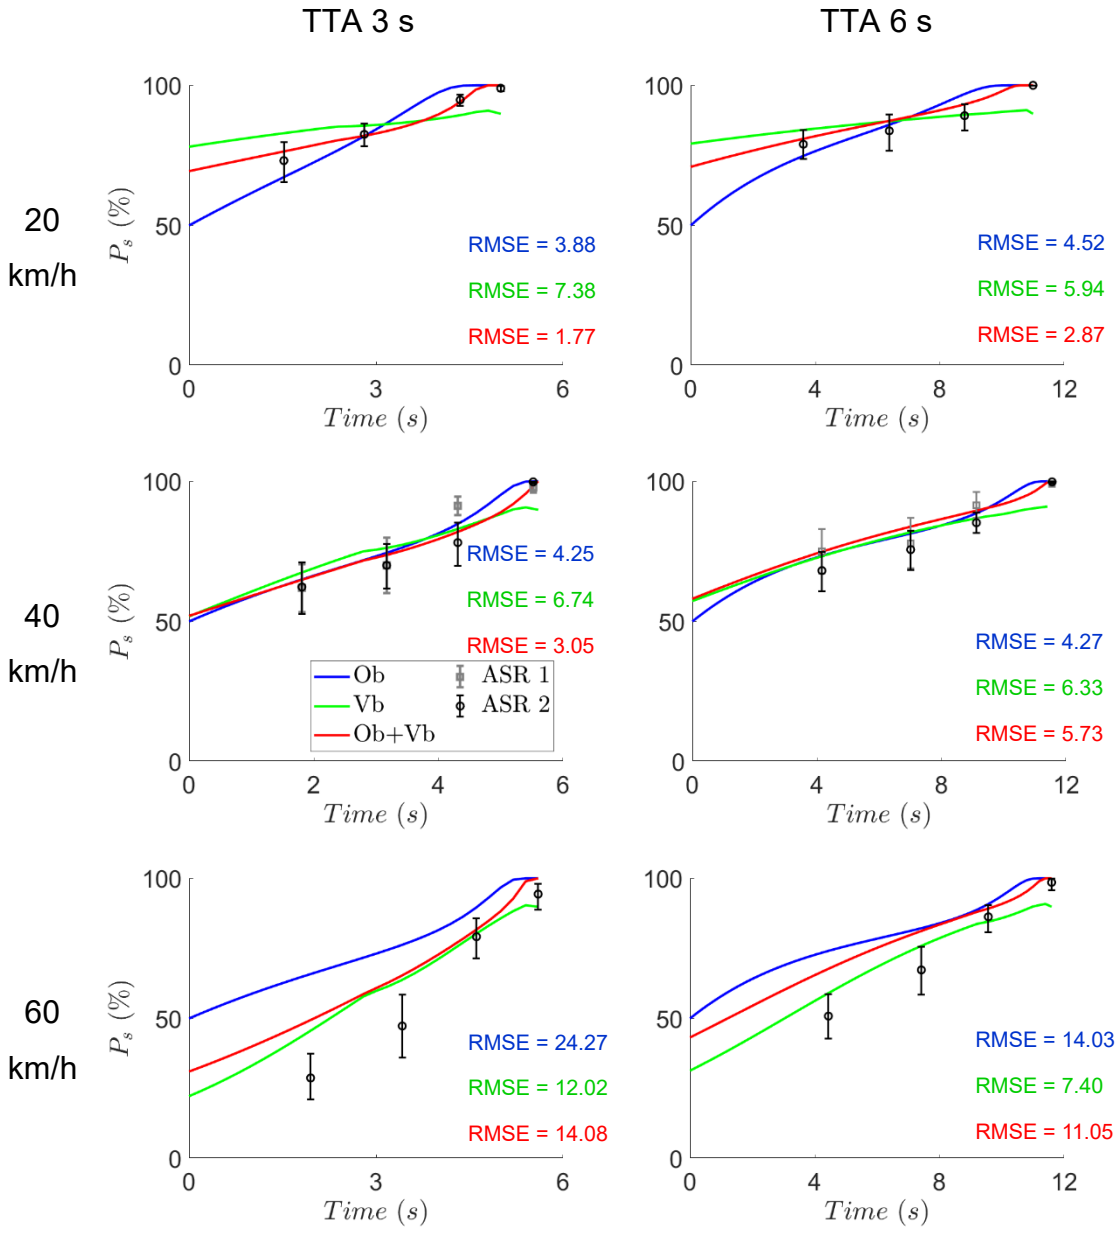

Figure 4.7 - Comparison between model predictions and average pedestrian beliefs for the deceleration manoeuvres. The RMSEs are regarding the new dataset only (i.e., Average Subjective Ratings 2)

In the deceleration scenarios (Figure 4.7), Vb and Ob+Vb generally demonstrated better performance compared to Ob in replicating pedestrian belief

patterns. While Ob's initial predictions failed to capture the influence of varying vehicle speeds and TTAs, its later predictions improved in accuracy as $\dot{\tau}$ values changed over time. Vb's initial predictions reflected the effects of speed and TTA, but its final predictions were less accurate due to the equalisation of the evidence of the two possible behaviours towards the end of each deceleration manoeuvres (Subsections 3.4.2 and 4.5.3). Ob+Vb's predictions effectively leveraged the early predictive strengths of Vb and the later strengths of Ob. Similarly to the constant speed scenarios (Figure 4.6), Ob+Vb's predictions in the 40km/h scenarios of Figure 4.7 support the validation of the suggested behaviour estimation mechanisms and the experimental paradigm, while Ob+Vb's predictions in the 20 and 60 km/h scenarios support its generalisability (especially for the lower speeds, in contrast to the constant speed scenarios).

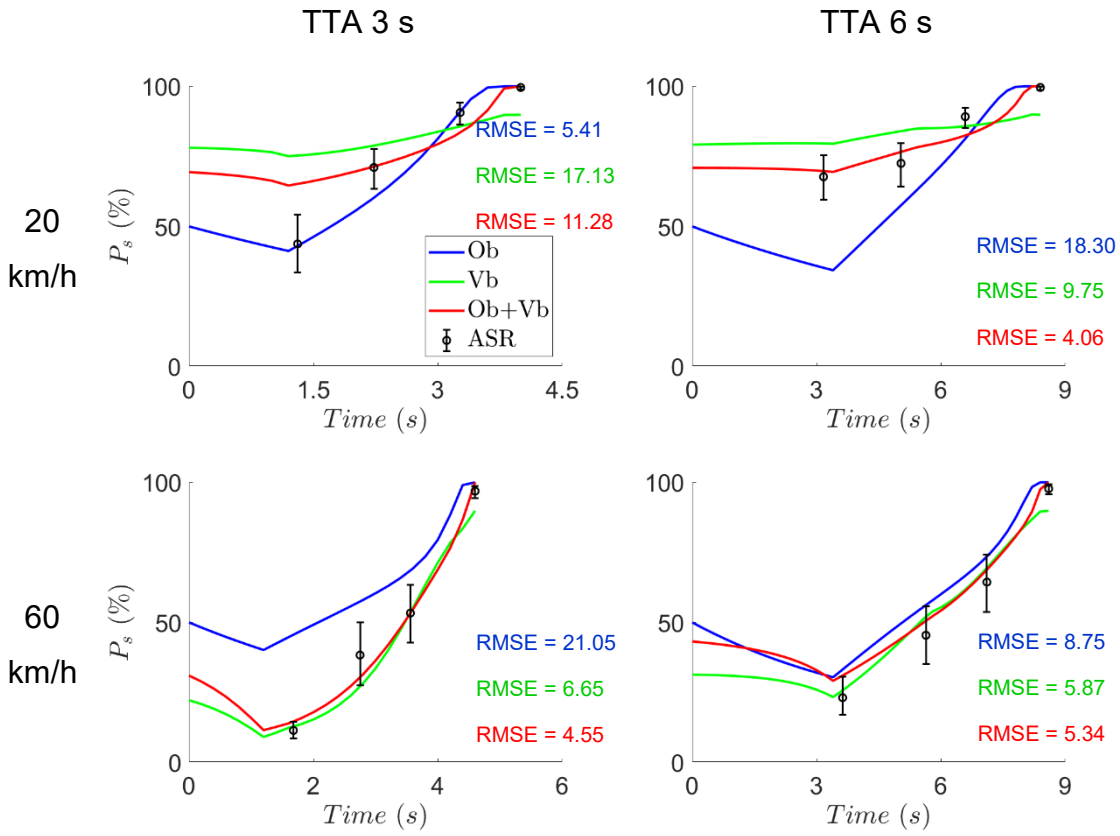

Figure 4.8 - Comparison between model predictions and average pedestrian beliefs for the two-stage deceleration manoeuvres

As can be seen in Figure 4.8, the two-stage deceleration scenarios were very similar to the mixed scenarios of the previous experiment. Even though kinematic details of the two-stage deceleration scenarios differed from the ones of the mixed scenarios, both the model predictions and the average pedestrians' beliefs followed similar patterns. In these scenarios Ob+Vb generally exhibited better

performance compared to Ob and Vb. The previously identified limitations of Ob's early predictions and Vb's later predictions, observed in the constant speed and deceleration scenarios respectively, were mitigated in the combined Ob+Vb model. The similarity in patterns between the two-stage deceleration scenarios and the mixed scenarios of the previous experiment provided a degree of validation for the models, particularly the Ob+Vb model. The fact that the model predictions, and especially the better performance of Ob+Vb, hold even when presented with slightly different kinematic profiles suggested a certain robustness and generalisability of the underlying model structures and the relationships they capture. This indicated that the models were not simply overfitting to the specific details of the original training data but were capturing more fundamental aspects of how pedestrians form and update their beliefs about the behaviour of approaching vehicles. The consistent mitigation of Ob's early and Vb's later prediction limitations in these new, two-stage deceleration scenarios further supported the validation and generalisability of the combined Ob+Vb model as a more comprehensive predictor of pedestrian beliefs across a range of dynamic situations.

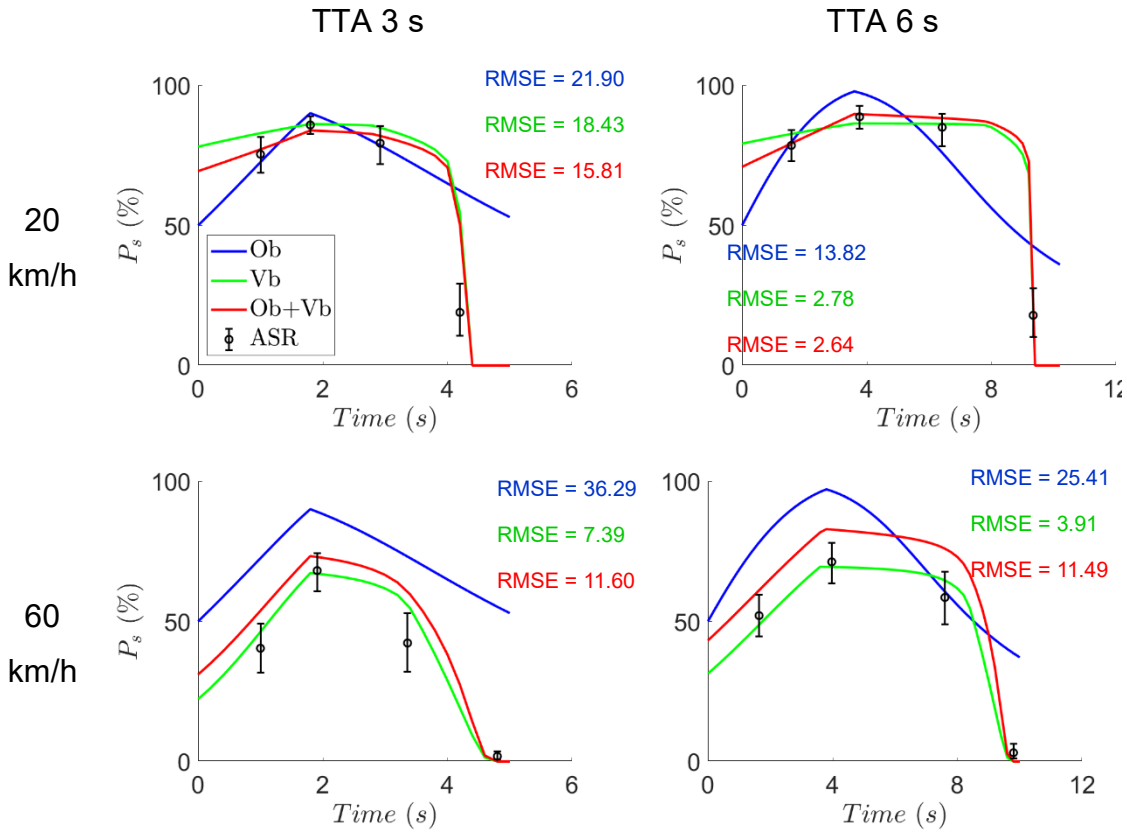

Figure 4.9 - Comparison between model predictions and average pedestrian beliefs for the short slowing manoeuvres

In the context of short slowing manoeuvres (Figure 4.9), Vb seemed to generally perform better than both Ob and Ob+Vb in predicting the average pedestrian beliefs regarding the approaching vehicles behaviour. The previously noted weakness of Ob when the vehicle speed is constant was again observed. This limitation was propagated to Ob+Vb due to its incorporation of observation-based evidence, leading to a slight overestimation of pedestrian beliefs that the vehicle was stopping (especially at higher vehicle speeds). Vb, conversely, exhibited the strongest agreement with average pedestrian beliefs during the short slowing manoeuvres. Importantly, the overall belief curve shape for these new short-slowing manoeuvres – an initial increase in belief of the vehicle stopping, followed by a decrease in that belief – is substantially different from belief curve shapes obtained in the previous experiment. The fact that Vb and Ob+Vb capture this general curve pattern reasonably well without any parameter refitting strongly suggests their ability to generalise to new scenarios and provides evidence that Ob+Vb was not overfitted to the previous dataset.

Across the majority of these 20 kinematic scenarios, when using the parameter settings that were obtained in Chapter 3, the Ob+Vb model provided the highest Spearman's rank correlation (along with Vb), the lowest overall RMSE and appeared (at least qualitatively) to be the most accurate in predicting the pedestrians' beliefs. These are strong indications that Ob+Vb was the best model (out of the three suggested ones) for capturing the average pedestrian beliefs regarding an approaching vehicle's behaviour. Furthermore, the consistent performance of Ob+Vb, even when applied to the new dataset and including novel scenarios (as discussed in relation to Figure 4.8 and Figure 4.9), suggested that the model possessed a notable ability to generalise and extrapolate to non-tested kinematic conditions of approaching vehicles. So, a key takeaway from the analysis so far is that the models fitted to the previous experiment demonstrated a considerable ability to predict pedestrian beliefs in this new dataset, which includes new scenarios. This supports the robustness of the underlying model structures and their ability to capture fundamental aspects of behaviour estimation.

The question of overfitting for a model as complex as Ob+Vb was a key consideration from Chapter 3. The results presented in this subsection, offered valuable insights into this question. The fact that Ob+Vb maintained its superior

performance (or at least remained among the top performers) on this new dataset, without being refitted, suggests that it was likely not substantially overfitted to the original dataset; overfitting would typically result in a more significant drop in its predictive accuracy when applied to new, unseen data.

The observation of a slight tendency of the average beliefs of the new experiment to be leaning a bit more towards the non-stopping behaviour in the four kinematic scenarios which were exactly the same in both experiments, indicated a difference in the pedestrians' prior beliefs between the two experiments. Hence, to further investigate the influence of prior beliefs within the model predictions in this new experimental context and to explore the potential for fine-tuning the model for this specific dataset, the next subsection focused on the refitting of the added margin parameter M. This would provide additional insights into the model's sensitivity and the extent to which its performance can be further optimised for the characteristics of the new experimental context.

4.5.2.2 Partial refitting (only parameter M)

This subsection explores the impact of refitting the added margin parameter (represented by M) on the predictive performance of the three models, for the new dataset. It is worth noting that fixing the values of the rest of the parameters to the best-fitted ones obtained in Chapter 3 and refitting only M, was a significantly less computationally expensive approach than fully refitting the models. Table 4.2 shows the best-fitting M values from Chapter 3 and the new optimal M values after refitting.

Table 4.2 - Parameter M values before and after refitting

	Old M	**Refitted M**
Ob	0	0
Vb	16.13	13.79
Ob+Vb	6.26	4.61

A quick comparison of the refitted and the previous M values (Table 4.2), shows that while M remained at 0 for Ob (consistent with findings in Chapter 3 where this value was optimal even when negative ranges were explored, indicating that there was not a model-driven bias towards the stopping behaviour when using the original search range, [0,100]), it was reduced for both Vb and

Ob+Vb. By comparing the first and second rows from Figure 4.5 it can be observed that the vertical displacement of the scatter points are consistent with the M value differences presented in Table 4.2. After refitting M, Vb and particularly Ob+Vb, once again offered better predictions of the pedestrian beliefs compared to Ob (Figure 4.5). The data points of Vb and Ob+Vb were translated downwards (only on the y-axis) – as a result of their reduced M values – while Ob's data points remained fixed since M remained constant and equal to zero. It is noteworthy that the reduction of the bias towards the stopping behaviour provided lower RMSEs for Vb and Ob+Vb, bringing them closer to the RMSE values obtained by the dataset of the previous experiment.

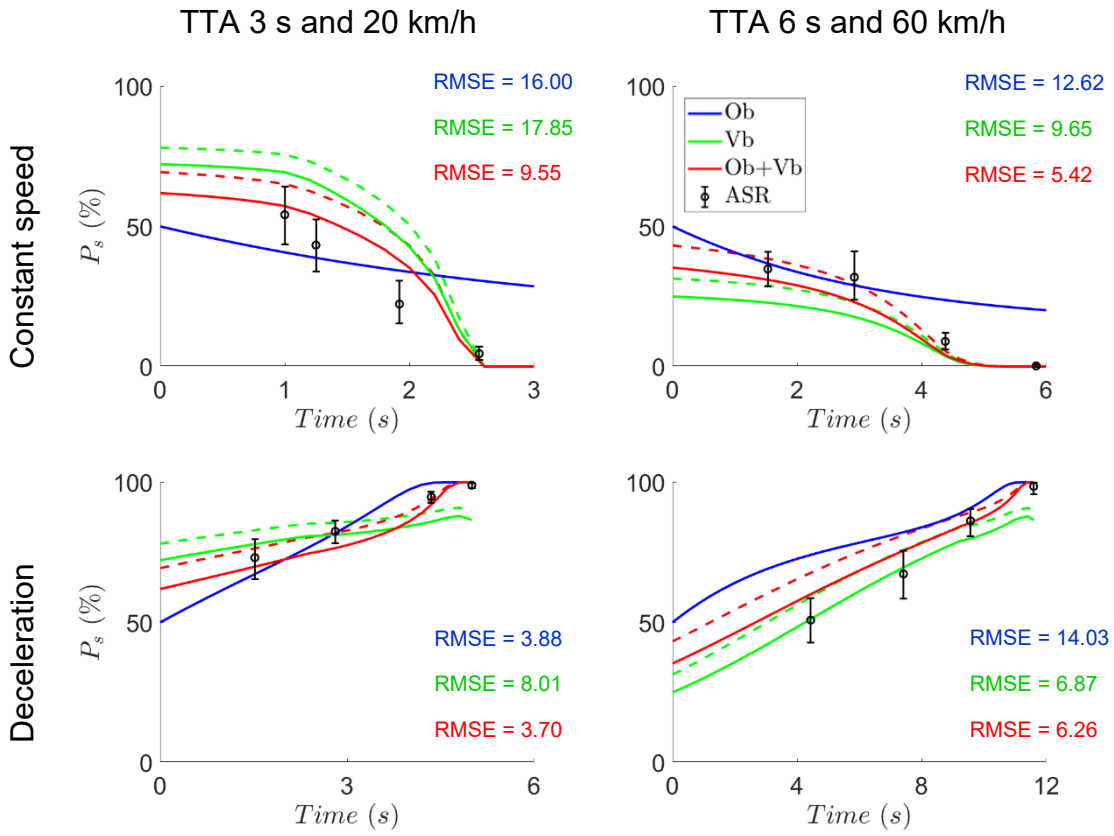

Figure 4.10 - Comparison between model predictions and average pedestrian beliefs for examples of constant speed and deceleration scenarios. The dashed curves illustrate the model predictions using the parameter settings of Chapter 3 and the solid curves the model predictions using the same parameter settings but with the refitted M. The RMSEs are regarding the model predictions using the refitted M only

Figure 4.10 illustrates the model predictions, after with and without refitting the parameter M, against the average pedestrian beliefs, in some of the constant speed and deceleration scenarios. Essentially, Vb's and Ob+Vb's prediction curves were translated downwards, while Ob's remained as before. This downward translation led to Vb and Ob+Vb capture the average pedestrian

beliefs even better (decreased the overall RMSE and maintained high rank correlation) than in Subsection 4.5.2.1, where the parameters settings of Chapter 3 were used.

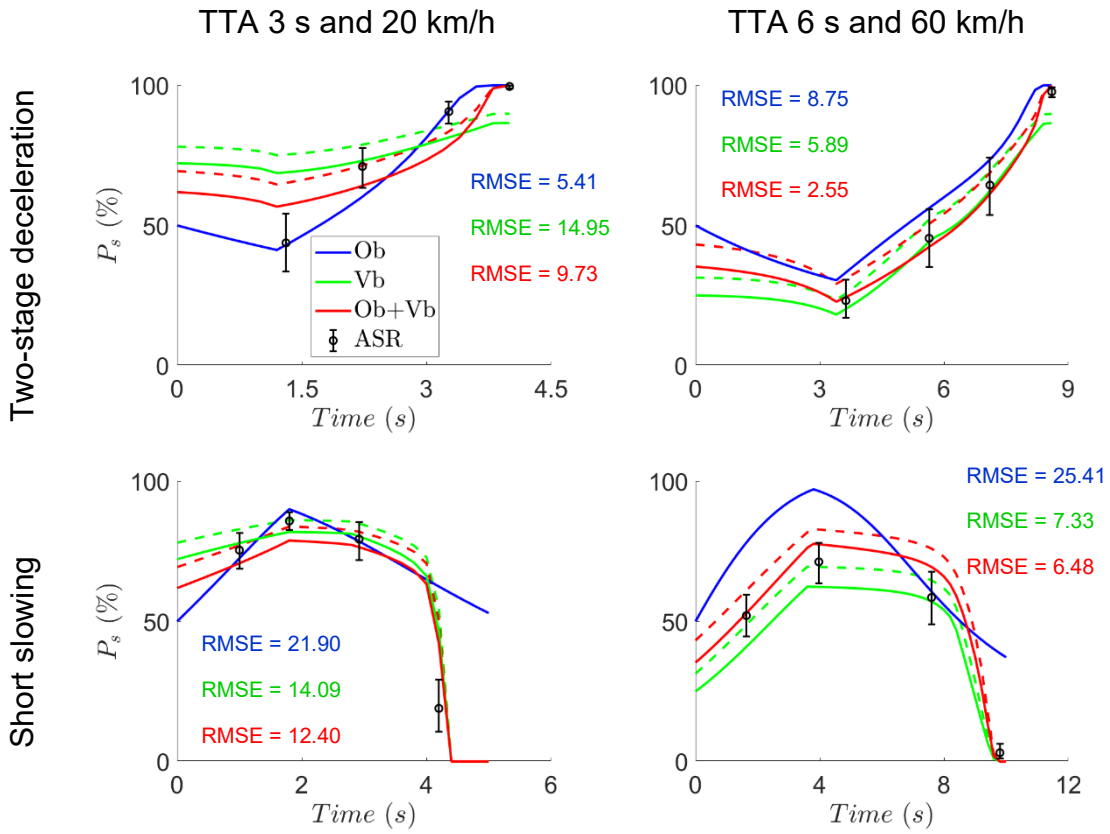

Figure 4.11 - Comparison between model predictions and average pedestrian beliefs for examples of two-stage deceleration and short slowing scenarios. The dashed curves illustrate the model predictions using the parameter settings of Chapter 3 and the solid curves the model predictions using the same parameter settings but with the refitted M. The RMSEs are regarding the model predictions using the refitted M only

Figure 4.11 demonstrates another example of the improved accuracy of Vb and Ob+Vb due to the downward translation caused by the smaller prior bias M. In the case of two-stage deceleration scenarios that shift seemed necessary especially for the first segments' judgements where the approaching vehicle was exhibiting almost-constant speed behaviour and the true P_S (participants' belief that the vehicle was stopping) was low. Regarding the short slowing scenarios, it can be seen that the belief overestimation of Ob+Vb observed in the previous subsection has been diminished due to the reduced bias towards the stopping behaviour of the vehicle. Thus, Ob+Vb seemed to be performing equally well with Vb now, in the short slowing manoeuvres.

In summary, the results of this subsection were obtained by maintaining the best-fitting parameter settings (all but M) of the previous chapter and then fitting the models to minimise the RMSE between the model predictions and the average pedestrian beliefs, through refitting the M parameter. Ob's predictions remained the same because the adjusted M value was the same as in the previous subsection, while Vb's and Ob+Vb's belief predictions were subtly translated downwards (closer to certain non-stopping behaviour belief or $P_s = 0$). Once more, Ob+Vb seemed to be performing better than Ob and Vb, based on both quantitative and qualitative comparisons.

4.5.2.3 Full refitting (all parameters)

Lastly, a rigorous full-model fitting sought to optimise the predictions of each model with respect to the new dataset. The panels in the third row of Figure 4.5 corroborate the resulting predictive accuracies of the three models after fitting their whole parameter settings on the new dataset. In consistence with both previous analyses based on different model fitting approaches (4.5.2.1 and 4.5.2.2), Ob+Vb demonstrated better predictive performance regarding pedestrian beliefs compared to Ob and Vb. The arguments behind this recurring theme are once again the decreased incidence of poorly predicted datapoints, high rank correlation and lower RMSE value.

A comparison of the new and the previous best-fitting parameter settings (Table 4.3), shows that:

For Ob, the best-fitting parameters remained similar, with the exception of μ_s, which was optimised at a higher value than before, meaning that there is a larger gap between the distributions of $P(\dot{\tau}|b_s)$ and $P(\dot{\tau}|b_{ns})$. The rest of the parameters were fitted to values of a similar level to the previous, indicating that Ob's best predictions are achieved by a specific parameter set, probably close to the refitted one.

For Vb, the full refitting resulted in more substantial changes in its optimal parameter values than Ob. The decrease in k_g suggested a reduced sensitivity to the evidence related to the driver reaching their goal, while the increase in k_{da} indicated a greater weight placed on the discomfort associated with deceleration. The substantial decrease in c_{pol} implied a diminished influence of the evidence related to the vehicle's pro-social behaviour. These changes collectively indicate

that the value-based considerations influencing pedestrian beliefs may have been recalibrated in the new experiment, with a greater emphasis on the immediate dynamics of vehicle deceleration and less on expected driver intentions. Furthermore, the reduction in M indicated a lower overall bias towards the belief that the vehicle would stop, in this new dataset, in line with the findings from the partial refitting in the previous subsection. Moreover, it is noteworthy that the product of $B \cdot M$ had a lower value after (~2.01) than before (~2.26) refitting, suggesting that the fixed bias towards the stopping behaviour (in absolute evidence terms) seemed to be lower in the new experiment compared to the respective bias of the previous experiment.

Ob+Vb also exhibited notable parameter adjustments upon the full parameter set refitting. Similar to the Vb model, k_g and c_{pol} decreased, suggesting a reduced influence of goal-related and pro-social evidence. Conversely to Vb though, k_{da} also decreased. As for the Ob part of the model μ_{ns}, μ_s, σ_{ns}, σ_s remained similar, suggesting that the stability shown in the refitting by the pure Ob model might have been transferred to Ob+Vb as well. Interestingly, the weighting parameter for the observation-based evidence β_O, increased, indicating a potentially stronger influence of perceptual cues derived from the vehicle's motion. Conversely, the weighting parameter for the value-based evidence β_V decreased, suggesting a relatively reduced contribution of inferred driver rewards and costs in the belief formation process for this new dataset. The time constant for forgetting past evidence T_f decreased substantially, suggesting a greater reliance on more recent observations, with the new value being closer to the T_f level of Ob. M decreased, mirroring the trend observed in the Vb model and indicating a reduced bias towards the stopping behaviour of the vehicle, once again consistent with the findings from the partial refitting in the previous subsection. Similarly to Vb, the product of $B \cdot M$ had a slightly lower value after (~1.17) than before (~1.25) refitting, suggesting that the fixed bias towards the stopping behaviour (in absolute evidence terms) seemed to be lower in the new experiment compared to the respective bias of the previous experiment. These changes in Ob+Vb suggested a recalibration of the integration process between observation and value-based evidence, with a slightly greater emphasis on immediate perceptual information, slightly less emphasis on the expected values

than in the previous experiment and a slightly smaller bias towards the stopping behaviour.

Table 4.3 - Best-fitting parameter settings

	Ob		**Vb**		**Ob+Vb**	
	Old	*New*	*Old*	*New*	*Old*	*New*
k_g	-		0.88	0.57	0.57	0.38
k_{da}			0.36	0.55	0.51	0.43
c_{pol}			0.83	0	1	0.49
μ_{ns}	-1	-1	-		-1	-0.90
μ_s	-0.47	-0.21			-0.75	-0.74
σ_{ns}	0.55	0.59			0.56	0.68
σ_s	0.57	0.54			0.58	0.68
T_f	4.56	5.41			100	1.97
β_O	-				0.32	0.60
β_V					0.53	0.37
B	0.17	0.12	0.14	0.17	0.20	0.45
M	0	0	16.13	11.82	6.26	2.60

Examples of the resulting model predictions due to the full refitting can be seen in Figure 4.12. Figures with model predictions, using the parameter settings obtained by the full refitting, for all 20 kinematic scenarios can be found in the Appendix B .

Figure 4.12 illustrates four examples (one for each driving manoeuvre) of model predictions versus average pedestrian beliefs comparisons. The initial TTA and initial speed of the vehicle is the same for all these four examples. The solid lines represent the model predictions obtained by using the fully refitted parameter settings, while the dashed lines represent the model predictions obtained by using the partially refitted parameter settings (only M) of the previous subsection. The patterns observed in the examples of Figure 4.12 were consistent with the patterns of the rest of the kinematic scenarios.

Regarding the constant speed scenarios, as can be seen in Figure 4.12, Ob's predictions decrease rate was larger because of the adjustment of μ_s to a higher value than before, meaning that the model introduced a bigger and clearer perceptual discrepancy between the two possible behaviours. That could possibly be explained by the fact that the new experiment had a bigger number of non-stopping behaviour segments than the previous experiment. Vb's predictions shifted upwards due to the slight increase of the product $B \cdot M$ (1.93→2.01). Lastly, Ob+Vb appeared to be averaging the adjustments of the two behaviour estimation mechanisms, making it have the best prediction accuracy of the average pedestrian beliefs.

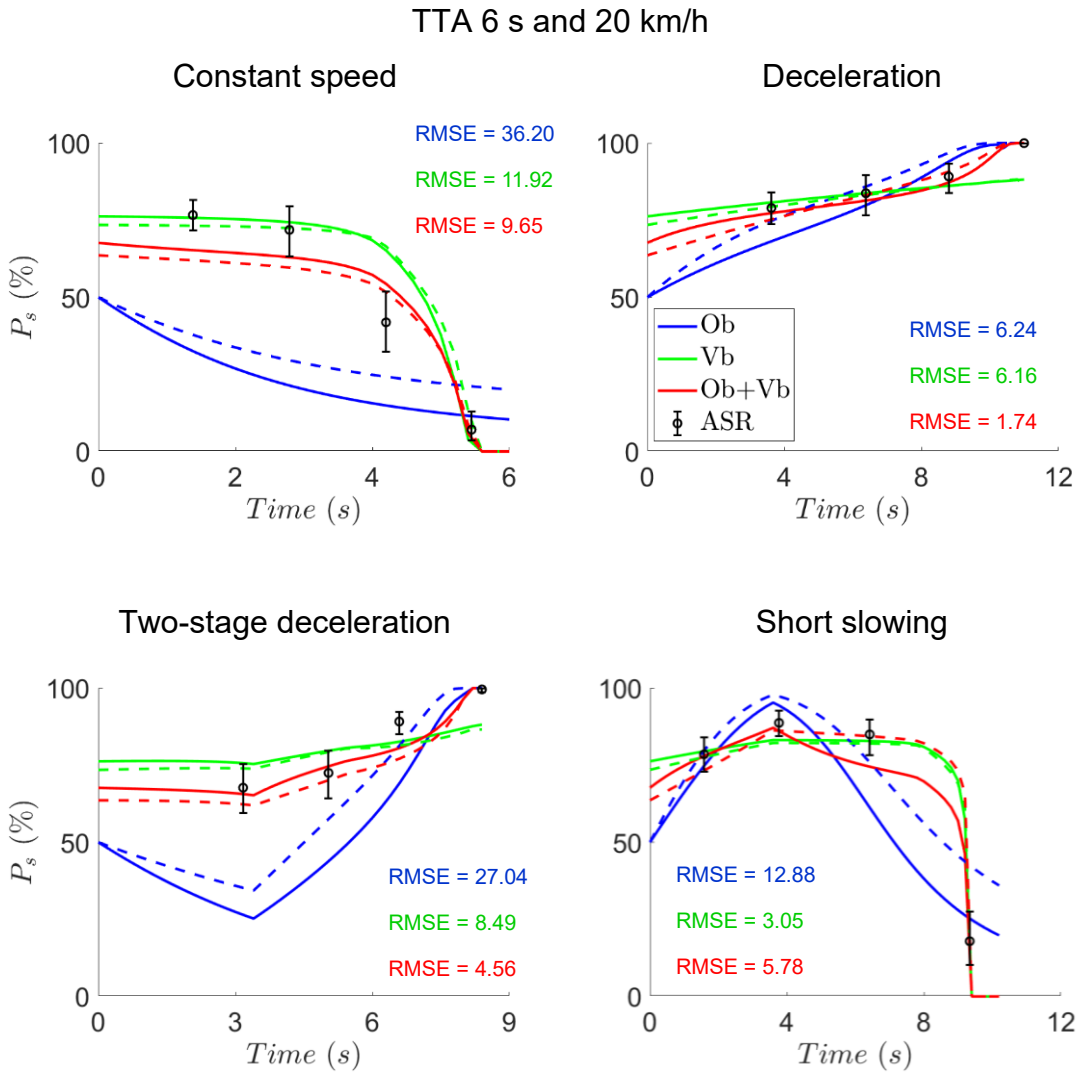

Figure 4.12 - Comparison between model predictions and average pedestrian beliefs for examples of all four driving manoeuvres. The dashed curves illustrate the model predictions using the parameter settings of Chapter 3 with the refitted M and the solid curves the model predictions using the fully refitted parameter settings. The RMSEs are regarding the model predictions using the fully refitted parameter settings only

Figure 4.12 demonstrates similar consistency with previously reported findings regarding the strengths and weaknesses of the models in the deceleration scenarios. As in the constant speed manoeuvres Ob's predictions rate was affected by the adjustment of μ_s. However, since the true vehicle behaviour in this case was the stopping behaviour, this adjustment of μ_s meant that Ob's predicted belief in the stopping behaviour now increased at a lower rate. Vb's predictions were once again shifted downwards due to the same reason as explained in the constant speed scenario of Figure 4.12. Ob+Vb averaged the abovementioned effects, achieving the best predictive accuracy.

The results for two-stage deceleration scenario presented in Figure 4.12, are again consistent with those previously reported when compared to the ones obtained from the different model fitting approaches of this chapter, but also when compared to the somewhat similar mixed scenarios of the previous experiment. From Figure 4.12, in the two-stage deceleration scenarios, Ob's predictions rate was affected depending on what was the true vehicle behaviour. More specifically, in the beginning of these scenarios when the vehicle was almost maintaining constant speed, Ob's predicted belief in the stopping behaviour decreased at a higher rate (consistent with what was observed in the constant speed scenarios), while in the later stages where the vehicle was decelerating, Ob's predicted belief in the stopping behaviour increased at a lower rate (consistent with what was observed in the deceleration scenarios). Vb's belief predictions were shifted downwards, as described before (this downwards shift was the same for all scenarios). The effect-averaging behaviour of Ob+Vb was once again apparent.

From Figure 4.12 and conversely to the two-stage deceleration scenario, early in the short slowing scenario, when the vehicle was decelerating, Ob's predictions regarding the stopping behaviour increased more slowly (mirroring the deceleration scenarios), whereas later, during the constant speed phase, its predicted belief in the stopping behaviour was decreased more rapidly (consistent with constant speed scenarios). Vb and Ob+Vb were affected the same way as explained before for all the other driving manoeuvre scenarios, after fitting their whole parameter settings to the new dataset.

Even though the last obtained parameter settings (refitting all parameters) obtained the lowest RMSE between model predictions and average pedestrian

beliefs, for all three models, the parameter settings with the refitted M was considered to be the most appropriate in balancing predictive accuracy with simplicity and generalisability to unseen data, possibly ensuring a candidate model that would perform well in real-world applications. That consideration was based on three key reasons. First, models with the refitted M parameter settings obtained RMSE values which were very close to RMSEs obtained by models with the fully refitted parameter settings. Second, it allowed a more direct and meaningful comparison of the models and their parameters across two experiments, rather than a completely new set of parameters. The examination of how the new dataset specifically influences the bias towards the stopping behaviour was facilitated, since all parameters except one remain fixed to their previously established optimal values. Lastly, the computational cost associated with refitting only a single parameter (M) is significantly lower than that of refitting a large number of parameters (7 for Ob, 5 for Vb and 12 for Ob+Vb).

Across all three parameter-fitting settings presented in the three previous subsections and across all 20 kinematic scenarios, the augmented model (Ob+Vb), consistently exhibited the highest Spearman's rank correlation, the lowest RMSE and appeared to provide the most accurate predictions of pedestrian beliefs both qualitatively and quantitatively. These consistent metrics strongly suggest that Ob+Vb's predictive performance is the best among the three proposed models. While this provides compelling evidence for Ob+Vb's effectiveness and suggests it is not overly complex or overfitted (given its strong performance on a new dataset with parameters fully or largely derived from a previous one), further investigation into the underlying behaviour evidence integration (Subsection 4.5.3) and a per-participant analysis (Subsection 4.5.5) will provide a more comprehensive understanding of the model's strengths and limitations. To complete this investigation and further understand these aspects, the Bootstrap Cross-Validation (BSCV) model selection technique (Subsection 4.5.4) was applied. BSCV offered a robust method for comparing model performance and generalisability on unseen data, accounting for overfitting, especially for the more complex Ob+Vb. Additionally, it allowed a more structured analysis of the obtained model parameters values.

4.5.3 Breaking down the evidence

This subsection provides a more in-depth analysis of the significance of evidence related to the approaching vehicle's behaviour. To reiterate, A_{ns} represents evidence of non-stopping behaviour, A_s represents evidence of stopping behaviour, A_{Ob} represents evidence derived from the observation-based (Ob) behaviour estimation mechanism and A_{Vb} represents evidence derived from the value-based (Vb) behaviour estimation mechanism (refer to Equations (3.15) and (3.16) for clarity). Essentially, this analysis provides insights of how the proposed models produce their predictions. However, since the constant speed, deceleration and two-stage deceleration scenarios of the new experiment are similar to the constant speed, deceleration and mixed scenarios, respectively, of the previous experiment, and the fact that another comprehensive analysis of evidence break down was presented in Chapter 3, this subsections' analysis will only address the short slowing scenarios.

Figure 4.13 illustrates the influence of varying initial speeds on the derived vehicle behaviour evidence. In short slowing scenarios with the same initial TTA, higher initial speed correlated with a shift toward initially believing that the vehicle will maintain speed rather than stopping. This is visible in Figure 4.13 as $A_s - A_{ns}$ being initially more negative for higher initial speeds. This happens in the model because higher speeds imply a greater reward for the driver/AV if the vehicle maintains its speed, leading to a stronger initial bias towards the non-stopping behaviour. This effect, which was apparent in Subsection 3.4.2 for all the other types of driving manoeuvres, has here been successfully propagated to short slowing scenarios as well.

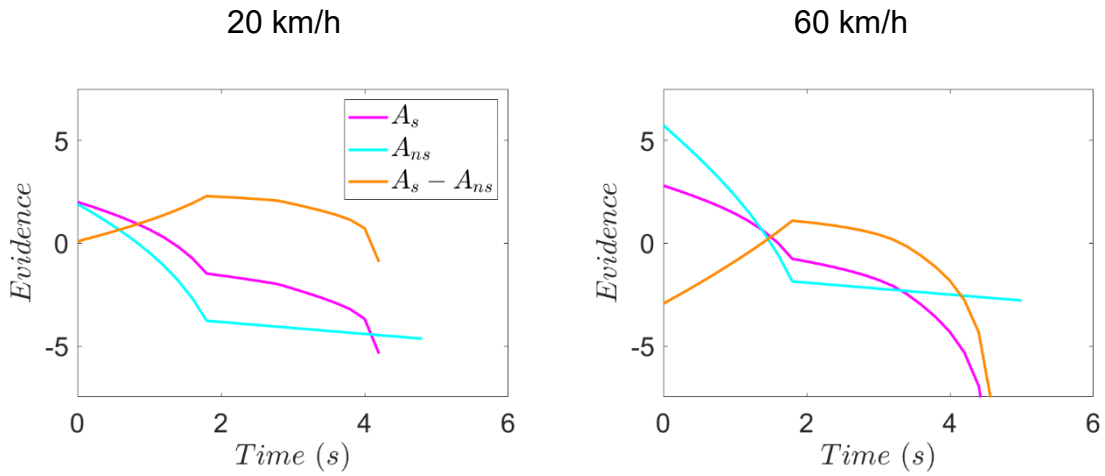

Figure 4.13 - Evidence of the two possible behaviours during short slowing manoeuvres. The initial TTA is 3 s in both panels

Figure 4.14, using a similar format, demonstrates the impact of different initial TTAs on vehicle behaviour evidence during short slowing scenarios. When the vehicle is initially further away (longer initial TTA), there is a shift in evidence towards the stopping behaviour rather than the non-stopping behaviour. This is visible in Figure 4.14 as $A_s - A_{ns}$ being initially larger for longer initial TTAs. The underlying logic in the model is that a larger initial distance translates to a lower necessary deceleration rate for the vehicle to stop comfortably and safely before the position of the pedestrian. This makes the act of decelerating more advantageous, resulting in a stronger initial tendency to predict the stopping behaviour. Importantly, the key observation from Figure 4.13 and Figure 4.14 is that the evidence difference $A_s - A_{ns}$ is the primary determinant of the shape of the Ob+Vb model's predicted beliefs, as observed also in Subsection 3.4.2 for the scenarios studied in the first experiment.

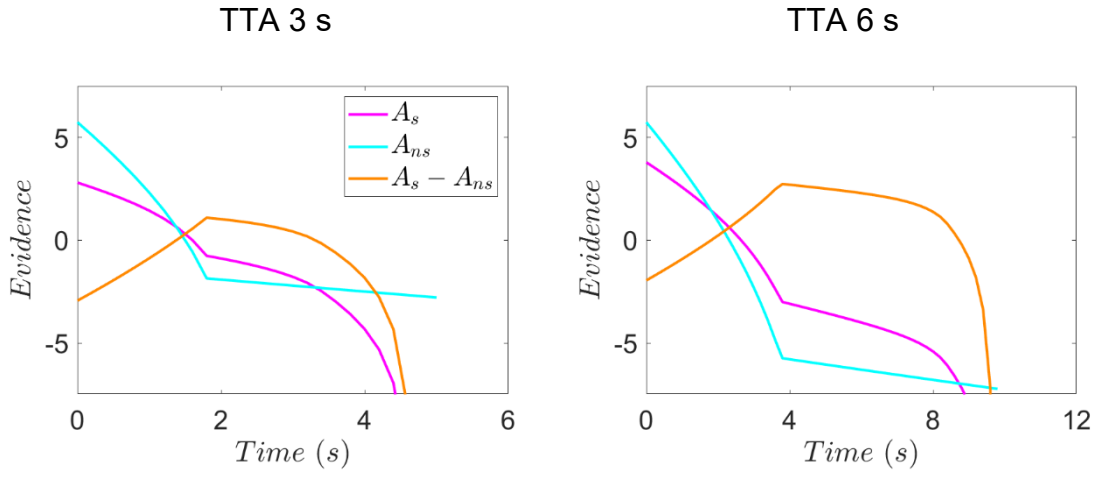

Figure 4.14 - Evidence of the two possible behaviours during short slowing manoeuvres. The initial speed is 60 km/h for both panels

Then the behaviour evidence are broken A_{ns} and A_s further down to the observation and value-based behaviour estimation components, as shown in Equations (3.15) and (3.16). To avoid any confusion the added soft-margin M has been omitted in Figure 4.15 and Figure 4.16, since it is a constant value and would be illustrated by the same horizontal line in all the following plots.

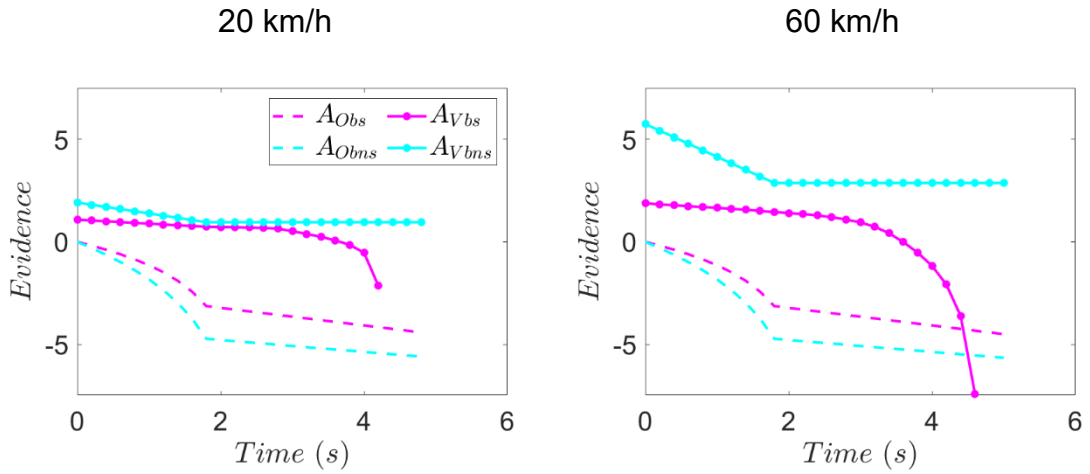

Figure 4.15 - Evidence of the two possible behaviours during short slowing manoeuvres, divided by behaviour estimation mechanism. The initial TTA is 3 s for all panels

Figure 4.15 depicts the influence of varying initial speeds during short slowing manoeuvres on the evidence computed by both the observation-based and value-based behaviour estimation components of the Ob+Vb model. Correspondingly, Figure 4.16 illustrates the effects of different initial TTAs during short slowing manoeuvres on the evidence derived from these same components.

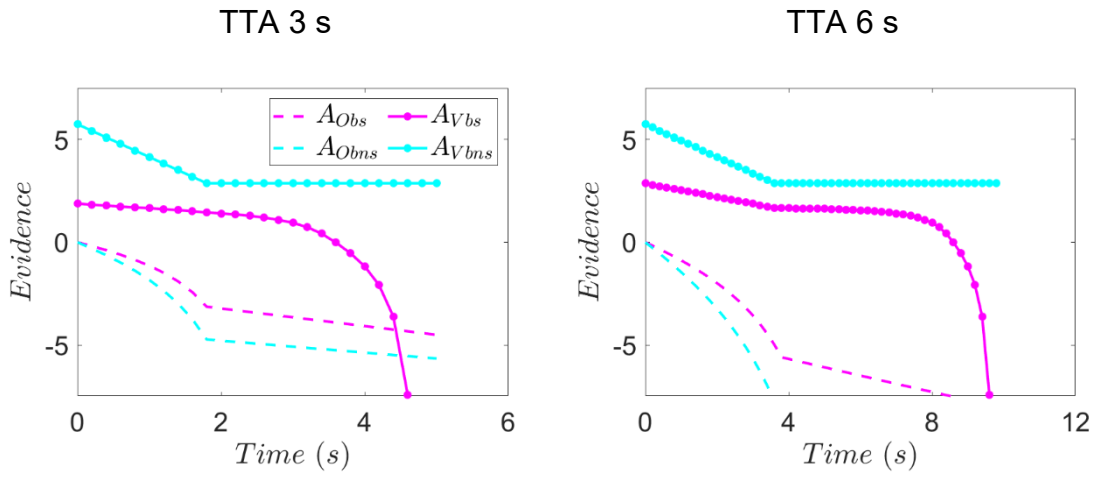

Figure 4.16 - Evidence of the two possible behaviours during short slowing manoeuvres, divided by behaviour estimation mechanism. The initial speed is 60 km/h for all panels

Figure 4.15 and Figure 4.16 reveal that initially, both A_{Obns} and A_{Obs} exhibit zero evidence values during short slowing, due to the lack of early $\dot{\tau}$ data, while A_{Vbns} and A_{Vbs} possess non-zero initial values. Specifically, $A_{Vbns}(t=0) > A_{Vbns}(t=0)$, which is consistent with Equations (3.4) and (3.5), given that $R_d(t) = R_c(t) - k_{da}a_{req}(t)^2 + c_{pol}$ (with c_{pol} insufficient to offset the deceleration discomfort cost). This confirms, once again, observations from the previous chapter, explaining why the observation-based approach struggles with early predictions while the value-based approach shows stronger predictive capability in the initial stages.

During the constant speed phase of short slowing scenarios (also during constant speed manoeuvres and the near-constant speed phase of two-stage deceleration scenarios), A_{Vbns} remains constant since the vehicle's speed is constant. Conversely, A_{Vbs} decreases towards $-\infty$ as the required deceleration to stop before the pedestrian increases with decreasing vehicle-pedestrian distance. Simultaneously, the evidence difference $A_{Obs} - A_{Obns}$ increases, as $\dot{\tau}$ observations continue to align with the actual vehicle behaviour (constant speed).

During the deceleration phase of the short slowing scenarios (also during deceleration manoeuvres and the decelerating phases of two-stage deceleration) A_{Vbns} decreases linearly towards 0, mirroring the vehicle's linear speed reduction. Similarly, A_{Vbs} decreases towards c_{pol} as both vehicle speed and required deceleration approach zero (ultimately converging at $A_{Vbs} - A_{Obns} \gtrapprox 0$). Concurrently, the evidence difference $A_{Obs} - A_{Obns}$ increases (reaching substantial values near the end) as $\dot{\tau}$ observations continue to reflect the actual vehicle behaviour, i.e., decelerating to stop. This dynamic explains the reduced accuracy of Vb's final-stage predictions compared to Ob's performance in predicting pedestrian beliefs at the conclusion of deceleration and two-stage deceleration scenarios.

This subsection has thus, once again, validated the combination of Ob and Vb in the Ob+Vb model by demonstrating the importance of the integration of different sources of behaviour evidence and highlighting the advantages offered by the two distinct behaviour estimation mechanisms, even in the context of a new approaching vehicle kinematic condition.

4.5.4 BSCV model selection

Besides testing the models on a new dataset, to ensure more robust model selection and generalisation, a bootstrap cross-validation (BSCV) analysis, as detailed in Sections 3.3 and 4.3, was implemented.

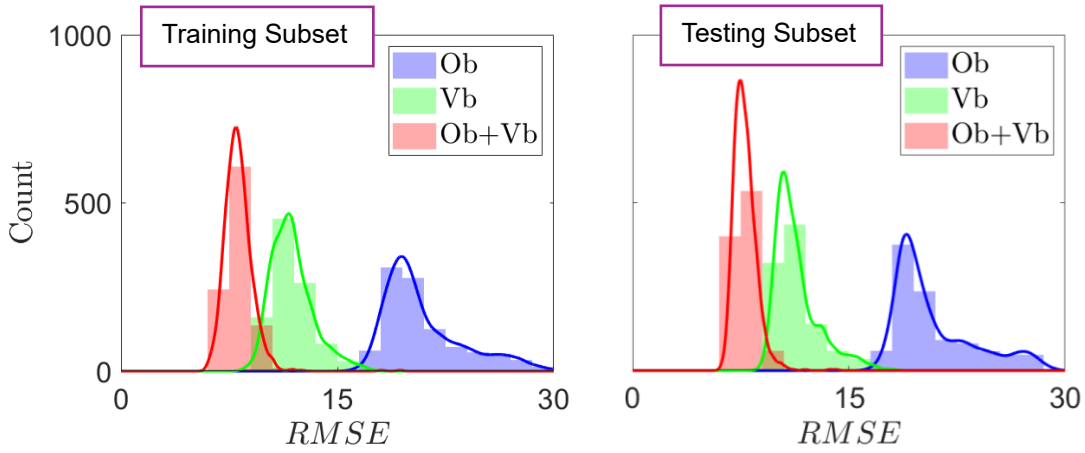

Figure 4.17 - Histograms (and their kernel-smoothed distributions) of the BSCV RMSEs of Ob, Vb and Ob+Vb

Figure 4.17 illustrates the distribution of the BSCV targeted metric, in this case the RMSEs, for the three proposed models. Consistent with previous observations, the Ob model exhibited the highest BSCV RMSEs, indicating inferior performance. A key finding from this analysis is that the BSCV results for the new experiment provide a clearer separation between the Vb and Ob+Vb models compared to the BSCV analysis presented in Subsection 3.4.3. As shown by their respective RMSE distributions in Figure 4.17, while there is still a slight overlap, Ob+Vb (red distribution) is more clearly shifted towards lower RMSE values than Vb (green distribution) and its peak is significantly lower. This suggests that Ob+Vb was the best-performing model. This clearer separation between the BSCV RMSE distributions of Vb and Ob+Vb in the new experiment might be due to the inclusion of more complex and varied kinematic scenarios (i.e., the 'Short Slowing' and 'Two-stage Deceleration' manoeuvres and a wider speed range), which possibly led to a more comprehensive evaluation of the models, allowing the integrating behaviour estimation mechanisms nature of Ob+Vb to demonstrate its better predictive performance, more clearly, on unseen data. Overall, these BSCV results are consistent with those from Chapter 3 regarding model selection but offer stronger evidence for Ob+Vb's better performance in this new and more comprehensive dataset.

Another detailed parameter investigation was conducted, in a similar manner to that in Subsection 3.4.4. Rather than employing an exhaustive grid search across a broad range of parameter values (as in Subsections 2.5.3), this analysis focused specifically on the parameter values derived from the bootstrap cross-validation (BSCV) procedure.

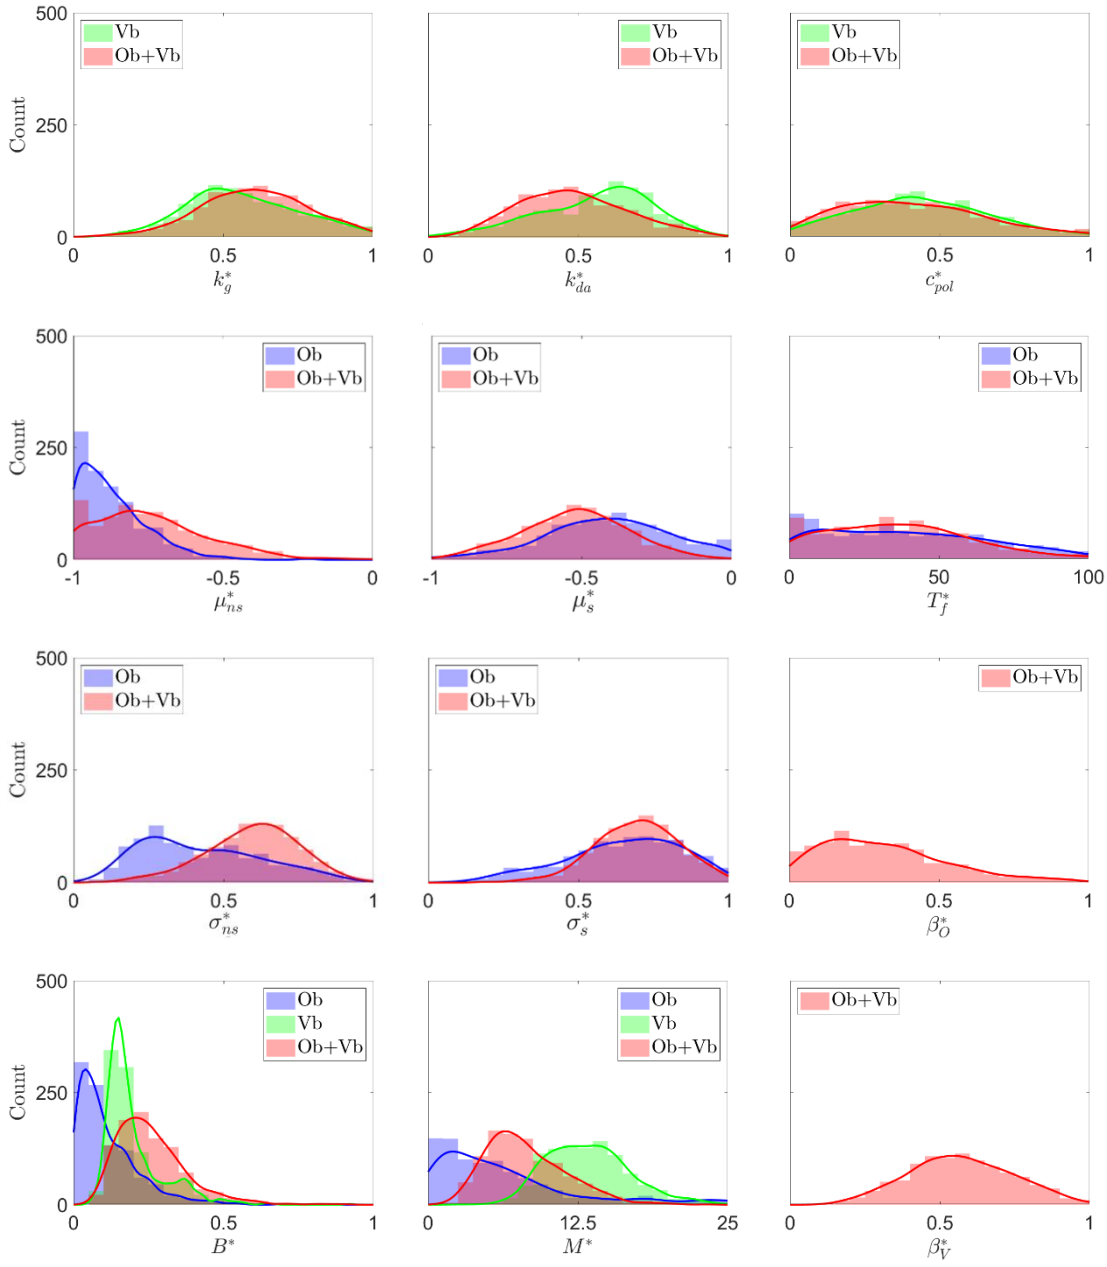

Figure 4.18 - Histograms (and their kernel-smoothed distributions) of the BSCV-obtained parameters of Ob, Vb and Ob+Vb

Figure 4.18 presents the histograms of these BSCV-obtained parameter values for each model, reflecting the range of values that yielded the best model fits across the BSCV resampled datasets. This BSCV parameter investigation conducted on the dataset obtained by the new experiment, produced parameter

distributions for the Ob, Vb, and Ob+Vb models that were remarkably consistent with those detailed in Chapter 3 (Figure 3.14). This replication across the two datasets obtained in the two experiments further reinforces the robustness and stability of the identified optimal parameter ranges for each model. The BSCV-obtained parameter settings and the indications of potential parameter redundancy, as illustrated in Figure 3.15, Figure 3.16 and Figure 3.17 in Chapter 3, which were mirrored in the current analysis (can be seen in the pairwise scatterplots in the Appendix D), appeared to not be specific to a particular dataset, possibly suggesting fundamental characteristics of how each model best captured pedestrians' beliefs. This consistency strengthens the confidence in the underlying mechanisms assumed by the models and the obtained optimal parameter values.

4.5.5 Per-participant model fitting

This subsection presents an analysis on the per-participant level, fitting the models to each participant's data from the second experiment. This analysis was motivated by important concerns regarding the simplifying assumption of using a single 'average pedestrian' parameterisation, as was done in all previous analyses.

The participants and humans in general are not identical. They vary in their cognitive processes, decision-making strategies, prior knowledge and many other factors. A single model fit to the aggregate data might hide these individual variations, essentially averaging them out and potentially misrepresenting the underlying processes for many participants (Ashby et al., 1994; Estes, 1956; Heathcote et al., 2000). Fitting models per participant allows capturing this heterogeneity and understanding how the model parameters vary across individuals, potentially leading to a personalised model that accounts for the specific characteristics and biases of each participant (Farrell and Lewandowsky, 2018).

Sometimes, even within a population, there might be distinct subgroups with different underlying processes (Lee and Newell, 2011; Unsworth et al., 2011; Vandekerckhove et al., 2008; Zhang and Luck, 2008). Fitting the model per participant could help identifying these subgroups by clustering individuals based on demographics and experience. This could reveal hidden patterns in the data that might have been missed by the previous whole-group-level analysis. Also,

fitting the model per participant could work as a check to whether the model assumptions hold for each individual or the majority of them. If the model does not fit well for some participants, it might be an indication that the framework/model is not specified correctly or that those individuals are using different strategies. In that case refinement of the model or considering alternative explanations could be motivated.

To explore the heterogeneity in how individual pedestrians' beliefs align with the models, a per-participant fitting procedure was performed. Specifically, for each of the 30 participants, the parameters B and M of all three models were individually optimised to minimise the RMSE between the model predictions and the participant's subjective beliefs across all 20 kinematic scenarios. These two parameters were chosen for the per-participant model fitting as they were directly related with the subjective judgments of vehicle behaviour in the experiment and hence were assumed to be capturing the most significant sources of inter-participant variability. Parameter B, the evidence scaling factor, might be reflecting the individual differences in response style; how a participant translated the available behaviour evidence into subjective ratings. Parameter M, the prior belief bias towards a certain vehicle behaviour, could possibly represent the individual tendency of each participant to believe that a vehicle will stop, likely influenced by personal experiences, risk tolerance and expectations. The remaining model parameters (i.e., those that control the calculation of the observation and value based evidence) were fixed to their respective best-fitting values obtained in the analysis of Chapter 3, as they were assumed to represent the more basic mechanics of the behaviour estimation cognitive process and to be more consistent across the population.

Figure 4.19 presents the distributions of the resulting RMSEs for each model across all participants. From the figure can be seen that Ob+Vb yielded the lowest overall prediction errors when fitted to individual participants, as evidenced by its RMSE distribution being concentrated at lower values compared to both the Ob and Vb models. This finding suggests that while model complexity could increase the risk of overfitting, the integration of both observation and value-based behaviour estimation mechanisms in the Ob+Vb model is important for accurately capturing the nuances of individual pedestrian beliefs regarding the behaviour of the approaching vehicle. The superior performance of Ob+Vb at the individual

level further strengthened its position as the most promising model among the three proposed.

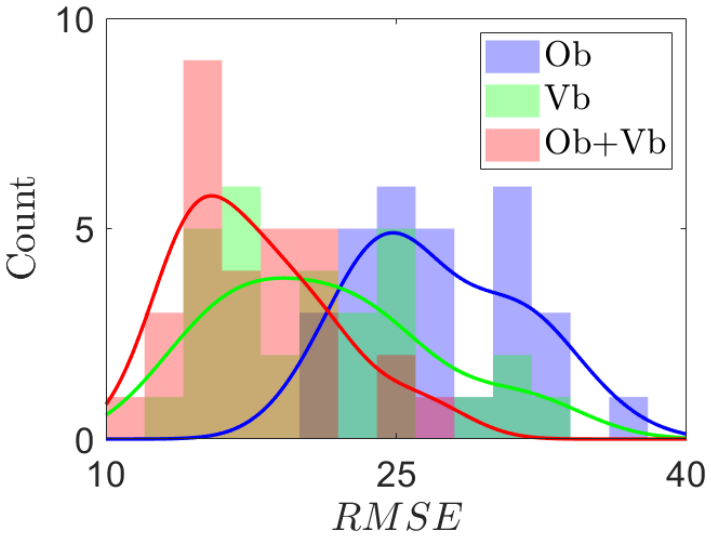

Figure 4.19 - Histograms (and their kernel-smoothed distributions) of the per-participant RMSEs of the three models

Notably, when comparing Ob+Vb's individual RMSE distribution (Figure 4.19) to the RMSE values of Ob+Vb obtained from fitting the model on the "average" pedestrian data (Figure 4.5), it appears that the RMSE values obtained from fitting the model to individual participants were larger than those obtained through the analysis on the aggregated data. This suggests that the predictions of Ob+Vb, even when its assumed individual-related parameters (B and M) are individually fitted, exhibit a greater average error when attempting to capture the nuances of individual pedestrian judgments compared to fitting the model to the averaged beliefs of a hypothetical "average" pedestrian.

This observation suggests that real human beliefs are likely affected by different individual factors and inherent variability in judgment, which are smoothed out when the models are analysed using the aggregated pedestrian beliefs. The RMSE values at the per-participant level generally seemed to be higher and more spread out than at the average level. This means that even though the model was successful at replicating the average pedestrian belief patterns, it was less accurate in predicting the beliefs of individual pedestrians. This could be due to variations in attention, risk perception, mechanical skills of their visual system, interpretation of cues, or even noise in their responses. The smoother RMSE distributions observed in BSCV analyses, which operated on averaged data, likely reflected the cancellation of some of this individual

variability, leading to a seemingly more accurate “average” prediction. To further investigate potential sources of individual variability and explore whether specific demographic or experiential factors were associated with differences in how well the Ob+Vb model captures pedestrian beliefs and the fitted model parameters, the 30 participants were grouped based on age: ≤50 vs. >50 years old, gender: Female vs. Male, and driving experience: Experienced (≥5 years of active driving experience) vs. Non/Novice (<5 years of active driving experience).

Figure 4.20 presents the grouped distributions of the Spearman's rank correlation (ρ_S) between the Ob+Vb model's predictions and each participant's beliefs for the abovementioned clusters. A comparison of these distributions across younger and older participants, female and male participants, and experienced drivers and non-drivers/novice drivers revealed no clear differences. This suggests that the degree of model's predicted beliefs rank ordering and individual participants’ subjective beliefs rank ordering alignment was not associated with age, gender, or prior driving experience.

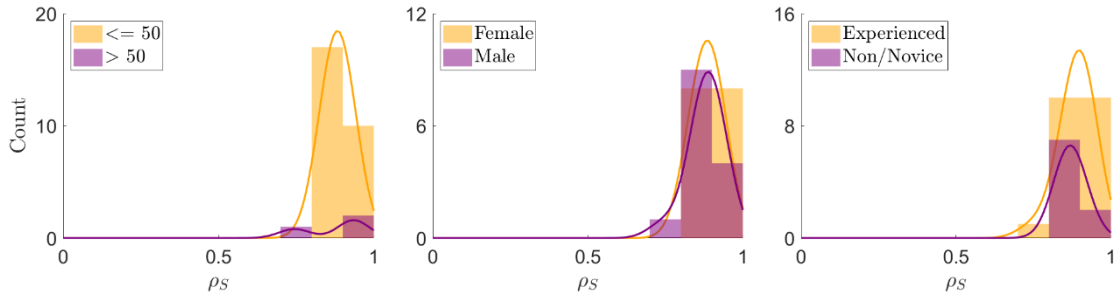

Figure 4.20 - Grouped distributions of Spearman's rank correlation between Ob+Vb predictions and individual pedestrian beliefs, by age, gender, and driving experience

Figure 4.21 displays the grouped distributions of the RMSE, providing a measure of the magnitude of the difference between the model's predictions and individual beliefs. A comparison of these distributions showed no difference in the model’s achieved RMSEs between female and male participants, nor between experienced drivers and non-drivers/novice drivers. This suggests that the model's predictive accuracy was not influenced by gender or prior driving experience. In contrast, age appeared to have a minor influence; the model’s achieved RMSEs of younger participants tended to be lower than those of older participants, suggesting that the absolute accuracy of the model was somewhat better for younger participants. In other words, younger participants’ beliefs were probably closer to the “average” pedestrian beliefs predicted by the model,

compared to older participants' beliefs – hence the slightly better predictive accuracy of the model for younger participants.

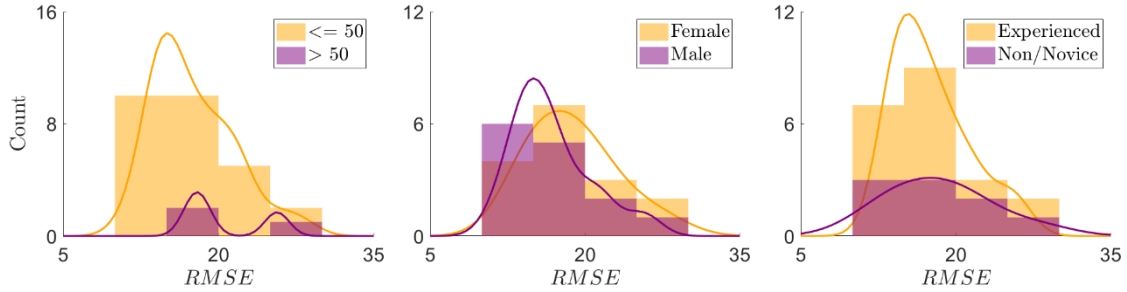

Figure 4.21 - Grouped distributions of RMSE between Ob+Vb predictions and individual pedestrian beliefs, by age, gender, and driving experience

Figure 4.22 presents the grouped distributions of the best-fitted B parameter. As a reminder, this parameter controlled the mapping from the model's calculated behaviour evidence (essentially the difference between the evidence for the vehicle stopping, A_s, and not stopping A_{ns}) to a subjective belief rating (essentially the subjective ratings provided by the participants in the experiment). More specifically, a higher B value implied a steeper translation of the predicted belief probability to the ratings provided through the questions of the behaviour estimation task (Appendix C). A comparison between the subgroup distributions of B revealed a noticeable difference due to age; B was greater for older participants than younger participants. However, this finding was likely affected by gender. Comparing B distributions for female and male participants revealed that there was a small subgroup of males for which the B parameter was fitted to its highest observed values; further inspection showed these were all male participants over 50. Therefore, the age effect that was suggested previously might have been an interaction between age and gender, though this is inconclusive due to the sample imbalance. The distributions of the B parameter for experienced drivers and non-drivers/novice drivers showed no noticeable difference in B, however non-drivers and novice drivers showed a slight tendency towards lower values of B. This suggests that prior experience with driving might have affected how confidently individuals translated the observed behaviour evidence into subjective beliefs about the vehicle's behaviour, with less experience in driving probably being associated with more conservative translation from behaviour evidence to subjective belief.

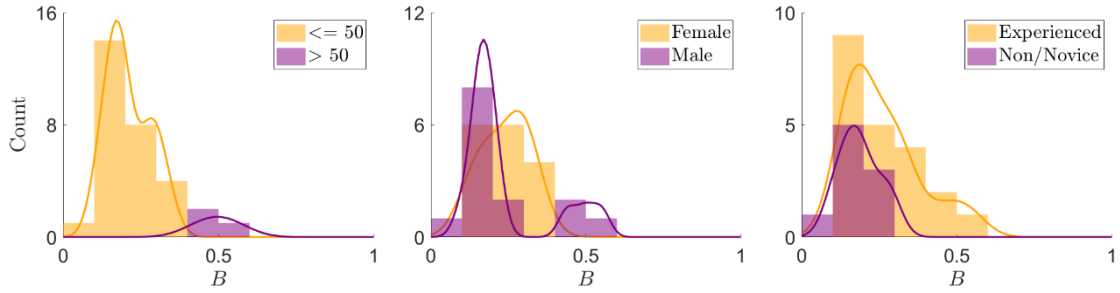

Figure 4.22 - Grouped distributions of the best-fitted translating factor (B) from behaviour evidence to belief probabilities for individual pedestrians, by age, gender, and driving experience

This variability in B across participants, as highlighted in Figure 4.22, underscores individual differences in how behaviour evidence is translated into subjective beliefs about the vehicle's behaviour. As illustrated in Figure 4.23, it is worth noting that for small absolute values of behaviour evidence, there is a substantial difference in the predicted beliefs of individuals with different B parameter values (red opening brace), while for big absolute values of behaviour evidence that difference is practically zero (red circle).

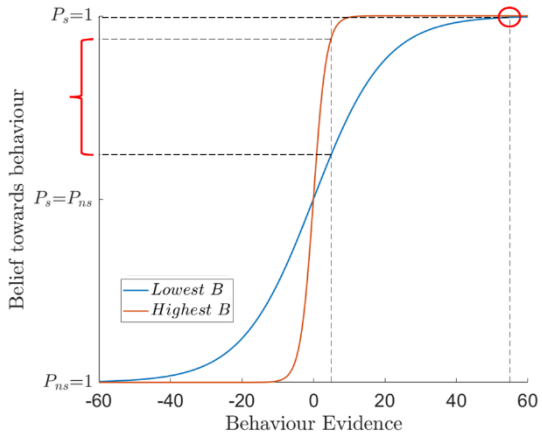

Figure 4.23 - Effect of the lowest and highest obtained scaling parameter B on the respective individuals' beliefs

Figure 4.24 illustrates the grouped distributions of the best-fitted M parameter, which represented the prior bias towards believing the approaching vehicle was stopping. A comparison of these distributions revealed that M was not affected by participants' gender or prior driving experience, as both showed no differences between their respective subgroups. On the contrary, a subtle tendency was observed for age; younger and older participants showed overall a similar achieved M, but with older participants' M values being clustered on the lower levels of prior bias. This indicated that older participants consistently held a smaller prior bias towards believing that the vehicle would stop.

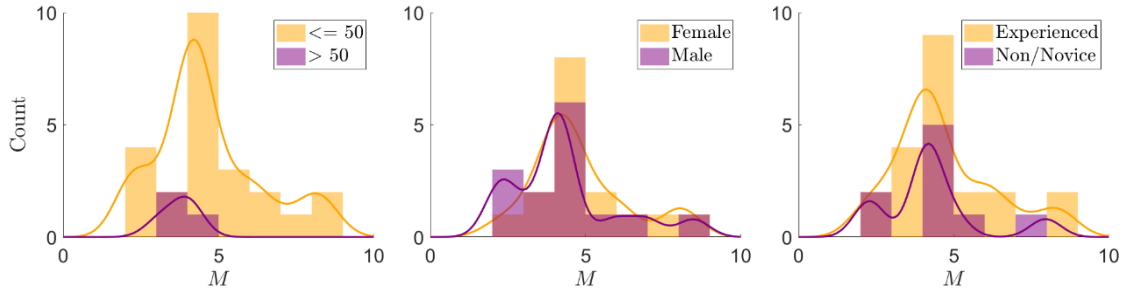

Figure 4.24 - Grouped distributions of the best-fitted prior bias towards yielding behaviour (M) for individual pedestrians, by age, gender, and driving experience

Figure 4.25 displays example scatterplots of Ob+Vb predictions versus observed subjective beliefs of four selected individual participants. These four selected participants were part of different demographic subgroups, i.e., a younger female, an older female, a younger male, and an older male. As discussed above, the obtained RMSEs of Ob+Vb at the individual level (Figure 4.19), were generally higher than when fitted to aggregated data. Figure 4.25 provides a visual illustration of this finding, showing higher prediction errors for individuals than the prediction error for the "average" pedestrian (Figure 4.5).

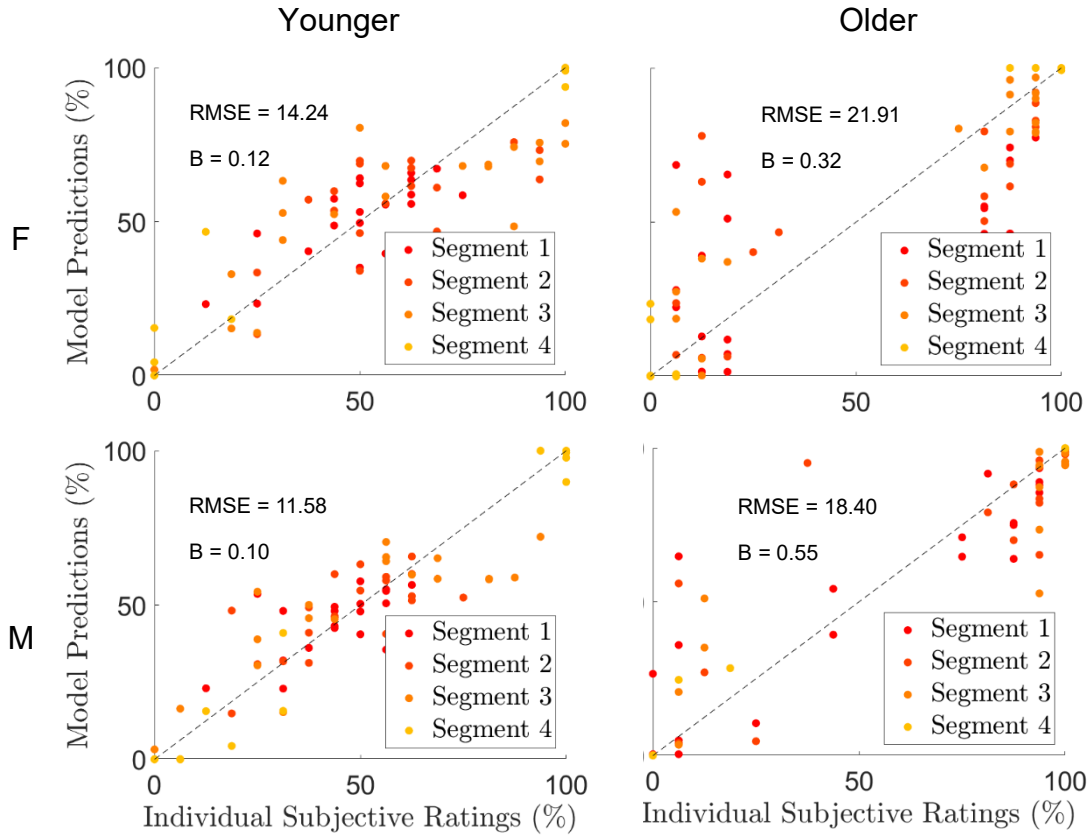

Figure 4.25 - Scatter plots of Ob+Vb predictions (y-axes) versus individual participants' beliefs (x-axes) for all 20 kinematic scenarios of the approaching vehicle. The panels provide illustrative examples of the model's fit for individual participants from different demographic subgroups, highlighting the variability in individual data and prediction error (RMSE)

Figure 4.25 offered a more granular view of the group-level results, within the age and gender groups, visually demonstrating the model's fits for four selected participants. These examples illustrate the kinds of individual variability that contribute to the per-participant RMSE distributions (Figure 4.19). For instance, the scatterplots for the two older participants (right column) exhibit more poorly predicted datapoints compared to the scatterplots for the two younger participants (left column). This is consistent with the trend noted earlier where the model's prediction error was slightly higher for older participants (Figure 4.21). Conversely, the model's accuracy did not seem to be affected by gender (comparison between rows). Lastly, as shown in Figure 4.25 by the individual B values, the pattern of the data points seemed to be influenced by B. In this instance, both the older example participants had a higher B value compared to the younger example participants, indicating a steeper relationship (as seen in Figure 4.23) between their ratings based on the presented evidence and the model's predictions across the different scenarios. This is evident from the fact that the data of both older participants were gathered close to the edges of the individual beliefs, while the data of both younger participants were spread across the whole belief scale, i.e., the older participants seemed to be more certain than the younger participants. However, that certainty did not necessarily mean that their beliefs were more accurate regarding the actual vehicle's behaviour.

4.6 Discussion

In the work presented in this chapter, the previously developed and suggested behaviour estimation models (Ob, Vb and Ob+Vb) were fitted to a newly acquired dataset. This new dataset was more comprehensive than the one from the previous experiment, as it had more tested kinematic scenarios (20 versus 18), a wider range of vehicle speeds (20-60 km/h versus 25-55 km/h), and included two new, more complex driving manoeuvres (Two-stage deceleration was a more realistic and better-designed version of Mixed manoeuvres, and Short slowing was not present in the first experiment at all). The reasoning behind designing another behaviour estimation experiment was to obtain a dataset through which the validation and generalisation of the suggested models would be enabled.

With the new experiment, empirical observations replicated patterns and trends that were observed in Chapter 2, validating the overall experimental design and empirical findings. The statistical analysis of pedestrian beliefs using a mixed-effects linear regression revealed significant main effects of initial vehicle speed (negative correlation with belief in stopping) and initial TTA (positive correlation), mirroring findings from the previous experiment. Furthermore, the manoeuvre type significantly influenced pedestrian beliefs, with constant speed leading to lower beliefs in stopping and deceleration to higher beliefs, relative to short slowing.

Model validation was achieved by replicating specific kinematic scenarios from the first experiment, namely, the constant speed and deceleration manoeuvres at 40 km/h initial speed (for both 3 and 6 s initial TTA). In these scenarios, Ob+Vb successfully predicted the average pedestrian beliefs with high Spearman's rank correlation and low RMSE, confirming the findings from Chapter 3. The overall belief patterns, such as the significant negative effect of vehicle speed and the positive effect of TTA, were also replicated, further validating the model's predicting behaviour. Another positive outcome was the consistency between the mixed scenarios of the original experiment and the two-stage deceleration scenarios of the new experiment, where Ob+Vb was able to capture the overall similarity and confirm the importance of behaviour-changing scenarios in the analysis of behaviour estimation.

Beyond validation through replication, the current chapter’s work also sought to investigate the generalisation capabilities of the models, particularly Ob+Vb. Generalisation refers to the model's ability to accurately predict pedestrian beliefs in novel situations not encountered during training, i.e., not tested previously (Bay and Yearick, 2024; Shepard, 1987). This was assessed by examining the model performance across a wider range of kinematic scenarios than those used in the original experiment. Specifically, the new dataset included extrapolated variations of initial vehicle speeds and new overall manoeuvres, not present in the earlier experiment. Perhaps the strongest evidence for the model's generalisability was indicated by its performance before any parameter refitting. Using the parameter settings derived from the first experiment, Ob+Vb still achieved a near-perfect Spearman's rank correlation and low RMSE when tested against the new dataset. The model's predictions for the new initial speeds (20 and 60 km/h) aligned well

with the trends observed for similar speeds in the first experiment (25 and 55 km/h), showing it could successfully extrapolate. The successful prediction of pedestrian beliefs in the new manoeuvres (two-stage deceleration and short slowing) provided more evidence for the model's ability to generalise beyond the specific conditions of the original training data, which could be considered a key step towards deploying such models in real-world applications. Furthermore, the inclusion of the more complex two-stage deceleration and short-slowing manoeuvres in the new dataset offered a more realistic test of the model's ability to capture pedestrian beliefs in more dynamic and potentially unpredictable traffic situations.

Analyses, mirroring the ones performed in Chapter 3, showed that all three models captured the general trends in pedestrian beliefs, though with varying degrees of accuracy. Similarly to Chapter 3, Ob exhibited the poorest predictive performance, when tested on the new dataset, primarily due to its inability to account for kinematic effects beyond deceleration ($\dot{\tau}$). Lacking access to speed and time-to-arrival (TTA), or equivalently, distance, Ob's predictive capacity was limited to scenarios involving a change in tau-dot. Consequently, it struggled to accurately predict beliefs in constant speed scenarios and during constant speed phases of two-stage deceleration and short slowing scenarios, producing similar predicted belief curves regardless of the specific conditions. In contrast, Vb addressed these shortcomings by incorporating reward functions related to speed (vehicle progress), TTA/distance, deceleration (deceleration discomfort), and pro-social behaviour (yielding politeness), as it did in Chapter 3. Thus, Vb captured the influence of a wider range of kinematic conditions on the pedestrians' beliefs. However, Vb's performance, when tested on the new dataset, once again appeared limited towards the end of vehicle approaches, particularly as the vehicle comes to a stop. This limitation arises from the near-equal evidence for both decelerating and constant speed behaviours in these situations. Given the importance of the difference in behaviour evidence (as discussed in Subsection 4.5.3), this near equality creates uncertainty for the model, even though the vehicle's stopping behaviour is often readily apparent to human observers. Ob+Vb ultimately combined the strengths of both Ob and Vb, successfully capturing the effects of all kinematic conditions and belief patterns present in the new dataset.

Quantitatively, Ob+Vb and Vb demonstrated near-perfect positive Spearman's rank correlations, while the Ob model exhibited a lower, though still relatively high, correlation. In terms of RMSE, Ob+Vb achieved the lowest value, followed by Vb and then Ob. These results were consistently obtained through all different parameter settings that were used (Subsection 4.5.2). The BSCV analysis further supported these findings, revealing that Ob+Vb consistently yielded the lowest bootstrapped RMSE values, indicating its superior predictive accuracy and robustness against overfitting. This result was more conclusive than the BSCV analysis in Chapter 3, where the BSCV RMSE distributions for Vb and Ob+Vb had a more significant overlap. The clearer discrimination in the BSCV analysis of the current chapter suggests that the design of the second experiment was more effective for model selection, allowing Ob+Vb's better performance to be more apparent. Furthermore, the parameter investigation using BSCV-obtained parameter ranges showed consistency with the findings of Chapter 3, indicating overall stability of the optimal parameterisations across different datasets on the average level.

The per-participant model fitting revealed that while Ob+Vb consistently showed the lowest prediction errors at the individual level, the RMSE values were generally higher than when fitted to the aggregated data, highlighting the inherent variability in individual pedestrian judgments. Analysis of individual differences based on age, gender, and driving experience showed some trends, particularly with older male participants exhibiting slightly higher RMSE and tendency to have a different mapping from behaviour evidence to subjective beliefs (higher B values) and a lower prior bias towards the stopping behaviour (lower M values).

Based on the consistent quantitative and qualitative superiority of Ob+Vb across both aggregated and individual levels, and supported by the model selection through BSCV, it seems appropriate to select Ob+Vb as the best of the suggested models and to recommend the parameter settings with the only refitted parameter being the prior bias (M) for use in subsequent studies or real-world applications, based on three key reasons. First, this approach is very efficient, as the computational cost of optimising only one parameter (M) is substantially lower than optimising multiple parameters (twelve in total for Ob+Vb) during model fitting. Second, its obtained RMSE value was very close to that of the full refitting. Third, this choice facilitated a more direct comparison of the models and their

parameters across the two experiments, rather than introducing a completely novel parameter configuration. More specifically, regarding M itself, the adjusted value was found to be lower than in the previous chapter. That could be explained by the design of the experiments. The original experiment included 72 segments, out of which the 30 were segments in which the vehicle maintained constant speed and 42 (56% of the total segments) in which the vehicle was decelerating. The new experiment included 80 segments, out of which 40 were segments in which the vehicle maintained constant speed and 40 (50% of the total segments) where it was decelerating. Therefore, the participants of the original experiment may have developed a small bias towards the yielding behaviour of the vehicle since they were experiencing it slightly more, and that could be detected as a larger M value.

5 General discussion

For safe and efficient road traffic interactions, pedestrians must process perceptual signals from approaching vehicles to interpret their behaviour and intentions. This behaviour interpretation challenge has become increasingly important with the emergence of automated vehicles (AVs). It might seem intuitive that pedestrians base their crossing decisions directly on perceived kinematic cues, such as the vehicle's current speed or distance. Yet, research suggests that behaviour estimations and road crossing decisions in traffic scenarios often rely on integrating information over time and expectations about the behaviour of other road users (Markkula et al., 2023; Tian, 2023). This thesis primarily investigated the cognitive mechanisms underlying how pedestrians estimate the behaviour (specifically, the intention to stop or not stop) of an approaching vehicle, specifically in the context of interactions with AVs. Understanding these mechanisms is important for improving road safety and designing AVs that can interact with pedestrians in a safe and intuitive manner. The current research addressed identified gaps in the literature, such as the lack of detailed empirical investigation into behaviour estimation from the pedestrians' perspective, the limited focus on vehicle manoeuvres that are observed in real-world beyond just traffic gaps between vehicles that travel with constant speed, and the absence of validated, generalisable computational models which explain the underlying behaviour estimation mechanisms.

The work presented in this thesis involved integrating principles from cognitive science, specifically Bayesian inference based on action observations and rational, value-maximising reasoning, into computational models. Two experiments using a CAVE-based pedestrian simulation environment were conducted to collect data on pedestrian beliefs under various kinematic conditions, including different vehicle speeds, Time-To-Arrival (TTA), and manoeuvre types (constant speed, constant deceleration, mixed/ two-stage deceleration, short slowing). Three main computational models were developed and evaluated: an Observation-based model (Ob) relying on perceived deceleration cues (specifically $\dot{\tau}$), a Value-based model (Vb) assuming pedestrians expect rational, value-maximising behaviour from the driver/AV, and an augmented model (Ob+Vb) combining both mechanisms.

This chapter presents the contributions of this thesis in detail. First, the key findings from the empirical and modelling work are summarised. Following, the theoretical and practical significance of these findings are explored, by contextualising the results in relation to the primacy of implicit kinematic cues and the specific cognitive mechanisms of behaviour estimation. Then, these mechanisms are discussed along with cognitive frameworks, pedestrian road crossing, and possible neurophysiological underpinnings. The final discussion sections are regarding the practical implications for AV design and road safety, the limitations of the current work and potential directions for future research.

5.1 Summary of key findings

The research yielded several novel findings regarding pedestrian belief updating. First, observation models based purely on low-level kinematic features like deceleration proved insufficient. The initial study (Chapter 2) demonstrated that while a Bayesian observer model (Ob) using $\dot{\tau}$ captured some of the pedestrians' belief updating, it failed in scenarios lacking clear deceleration, such as constant speed approaches where initial speed and TTA clearly influenced beliefs. This highlighted the limitations of relying solely on perceptual processing in realistic traffic interactions.

Second, the accuracy of pedestrian belief predictions was improved by incorporating expectations through value-based reasoning. The value-based model (Vb) introduced in Chapter 3, which modelled driver rewards and costs (speed, comfort, politeness), was able to capture the influence of initial speed and TTA, addressing a major limitation of Ob. However, Vb was less accurate in predicting pedestrians' beliefs when the vehicle was very close to coming to a stop, where the calculated utilities for the stopping behaviour versus the non-stopping behaviour became less distinct.

Third, the augmented model combining both mechanisms (Ob+Vb) consistently provided the most accurate predictions of pedestrian beliefs across both experiments (Chapters 3 and 4). It integrated the strengths of Ob (i.e., accuracy with clear deceleration cues, especially later in the vehicle's approach) and Vb (i.e., capturing initial expectations based on speed/TTA). This combined

model achieved near-perfect correlations and the lowest prediction errors (RMSE) against average pedestrian belief data.

Fourth, Ob+Vb demonstrated robustness and generalisability (Chapter 4). It replicated Chapter 3's findings when tested on a new participant group and successfully predicted beliefs in novel scenarios involving different speeds and new manoeuvre types (‘Two-stage deceleration', 'Short slowing'). This suggests that Ob+Vb captured underlying cognitive mechanisms rather than just fitting specific data patterns.

Fifth, both experiments confirmed that pedestrian beliefs are significantly influenced by initial vehicle speed (higher speed correlating with lower belief in stopping), initial TTA (longer TTA correlating with higher belief in stopping), and the specific driving manoeuvre and its history. Early and clear deceleration was particularly effective in signalling yielding intent.

Finally, while models were primarily fitted to average data, per-participant fitting (Chapter 4) confirmed Ob+Vb's better prediction performance at the individual level, albeit with higher prediction errors than when fitted to data averaged across participants. The results also suggested potential links between age and model parameters governing belief translation (B) and prior bias (M), indicating individual heterogeneity in belief interpretation and reporting.

5.2 The primacy of implicit kinematic cues

A main theme of this thesis was the fundamental role of implicit communication through vehicle motion. As stated in Chapter 1, researchers have suggested that such kinematic cues are important for pedestrians when they are deciding to cross or not the road (Dey and Terken, 2017; Lee et al., 2021; Rasouli et al., 2018). The work presented in this thesis indicated that these kinematic cues are important for the mechanisms with which pedestrians estimate an approaching vehicle’s behaviour, as well. This emphasis on kinematic interpretation can be found in several key studies. For instance, the work of Tian et al. (2023), reporting on the same experiment as presented in Chapter 2 of this thesis, documented how parameters like vehicle speed, TTA, manoeuvre type influence both pedestrian crossing decisions and their subjective estimations of

vehicle behaviour. More specifically, this thesis and the abovementioned work indicated that the visual cue $\dot{\tau}$ is a valuable indicator for the stopping behaviour of a vehicle and that there could possibly be a relationship between crossing decisions and pedestrians' judgements, suggesting the use of behaviour estimation mechanisms during road crossing decisions by pedestrians. While these findings are shared between the two works, this thesis' proposed model (Ob+Vb) extends them by offering a detailed mechanistic explanation of how different sources of related information are processed and dynamically integrated into beliefs. Ackermann et al. (2019) highlighted vehicle deceleration as a key communication signal. They also found that higher deceleration rates and lower speeds facilitate quicker detection of vehicle deceleration by pedestrians, which aligns with the findings of this thesis, as higher deceleration rates and lower speeds led to earlier pedestrian beliefs that the vehicle was stopping. The detailed investigation of different deceleration profiles in this thesis extended the findings of Ackermann et al. (2019) by showing how not only the rate but also the timing and consistency of deceleration influence beliefs dynamically over the entire vehicle approach. The current thesis moved beyond mere deceleration detection also by modelling the continuous belief-updating process and incorporating not only visual cues like $\dot{\tau}$, a concept also discussed by Tian et al. (2023) and Ackermann et al. (2019) in relation to the "Tau-Hypothesis", but also speed and TTA-related cues (required deceleration to come to a stop).

The prioritisation of vehicle kinematics by the pedestrians has been supported in the literature. Studies have shown that when explicit eHMI messages are not aligned with the vehicle's actual movement, pedestrians default to relying on the vehicle kinematics (Dey et al., 2021; Rezwana and Lownes, 2024) and that vehicle behaviour (yielding versus non-yielding) was the primary determinant of crossing willingness, regardless of the pedestrian's knowledge of the vehicle's automation status or its external appearance (Dey et al., 2019). Moreover, this continuous processing of the vehicle's kinematic cues in pedestrians' crossing decisions, has been supported by studies using multimodal data (e.g., including eye-tracking; Lyu et al., 2024) and also from the driver's perspective, where pedestrians' decisions were affected by different braking strategies (Yang et al., 2024). These findings align with this thesis' argument that kinematic interpretation

is the basis of the pedestrians' road crossing decisions and consequently of their understanding of the traffic interaction.

While this thesis focused exclusively on implicit kinematic cues, it is interesting to consider these findings also in the context of explicit communication, particularly eHMIs. Researchers who have compared implicit and explicit communication, have suggested a complex relationship between the two. eHMIs could enhance perceived safety, trust, and willingness to cross (Dey et al., 2019), but their effectiveness is context-dependent and requires learning and familiarity (Lee et al., 2022). Based on the argument stated before that kinematic interpretation forms the baseline for pedestrians' understanding, explicit eHMIs could act as a supplementary layer, which would be useful for resolving ambiguity in situations when kinematic cues are inherently weak, e.g., for low speeds, short distances, unclear priority (Lau et al., 2022). Thus, this supplementary role for explicit signals, specifically eHMIs, indicates the primary importance of the vehicle's kinematic cues in pedestrians' understanding of the vehicle's intentions (Clamann et al., 2017; Lee et al., 2021). This further validates this thesis' main focus on modelling the interpretation of implicit kinematic cues.

Past research on implicit communication has often focused primarily on the role of deceleration on both pedestrians' road crossing decisions and pedestrians' inferences about the vehicle's intent (Ackermann et al., 2019; Dietrich et al., 2020; Lee et al., 2022; Petzoldt et al., 2018; Tian et al., 2023). The present research built on the importance of deceleration on the pedestrian beliefs regarding the vehicle's behaviour. However, the findings indicated that pedestrians possibly integrate and process other kinematic information besides its deceleration, as well (e.g., vehicle speed, distance, TTA). A contribution of the current thesis was not only to confirm the importance of kinematic cues but to also offer and validate a computational model (Ob+Vb) that could explain the cognitive mechanisms with which pedestrians possibly integrate these implicit signals with their internal expectations to form and update their beliefs about the behaviour (stopping or not) of an approaching vehicle.

5.3 Understanding the behaviour estimation mechanisms

5.3.1 Behaviour estimation mechanisms and their integration

The work of this thesis provided significant support to the hypothesis that human pedestrians employ an integration of perceptual processing and expectation-based reasoning when interpreting the behaviour of approaching vehicles. The limitations of the purely observation-based model (Ob) and the purely value-based model (Vb), contrasted with the success of the combined Ob+Vb model, suggested that neither mechanism alone is sufficient to fully capture pedestrians' beliefs, and the underlying cognitive process of behaviour estimation involved.

This integration aligned with broader theories in cognitive science that suggest interaction between perception and cognition in action understanding, situation awareness, and decision-making under uncertainty (Markkula et al., 2023). In particular, this integration was formalised by combining the likelihood of sensory evidence given a hypothesis – similar to the Ob's processing of observed kinematic cues (Baker et al., 2009; Dindo et al., 2011; Pezzulo et al., 2013; Vilares and Kording, 2011) with the prior beliefs or expectations about the hypothesis – similar to Vb's assumptions about agent rationality (Jara-Ettinger et al., 2020; Lucas et al., 2014; Markkula et al., 2023; Wright and Leyton-Brown, 2017). The human mind continuously goes beyond the raw data of experience, making inductive inferences in uncertain environments, a process well-described by Bayesian principles (Griffiths et al., 2008; Vilares and Kording, 2011). The dynamic and complex nature of road crossing interactions, which are characterised by uncertainty and the need to anticipate others' actions, requires such an integration; static perceptual cues or fixed expectations are inadequate given the dynamic situation (Markkula et al., 2023; Wang et al., 2025). The success of Ob+Vb demonstrated the applicability of these cognitive principles, often explored in controlled laboratory settings (Baker et al., 2005, 2017, 2009; Dindo et al., 2011; Jara-Ettinger et al., 2020; Pezzulo et al., 2013; Zhi-Xuan et al., 2020), to the challenging, applied domain of pedestrian-vehicle interactions. While each of the two cognitive mechanisms described above, have been successful on their own, modelling complex behaviour estimation requires

considering both the likelihood of observations given intentions and reasoning about the other agent's goals and expected rationality.

The necessity of both the Ob and Vb components for optimal model performance implied that pedestrian behaviour estimation likely involves more than a simple summation of perceptual evidence and prior expectations. While the Ob+Vb formalised this integration with static weights, its success across such varied conditions suggested that the underlying cognitive process likely involves a dynamic relationship between these sources of evidence. This relationship was supported by the complementary performance of the two component models, where each was effective in areas where the other was limited. Pedestrians might rely more on their expectations about rational driver behaviour (the Vb component) when perceptual cues of deceleration are weak or ambiguous, such as when the vehicle is distant or maintaining a constant speed. Conversely, as clear kinematic evidence emerges (e.g., noticeable deceleration), the behaviour evidence might shift towards the observation-based component (Ob). This dynamic adjustment would explain why the combined model succeeds across a wider range of scenarios than either component alone. Such dynamic integration is characteristic of adaptive behaviour in complex systems (Markkula et al., 2023; Wang et al., 2025).

The value-based component (Vb) could be interpreted through the lens of computational rationality, affordance theory, and optimal control principles (Gawthrop et al., 2011; Hoogendoorn and Bovy, 2003; Jara-Ettinger et al., 2020; Lio et al., 2020; Markkula et al., 2023; Wang et al., 2015). Pedestrians appear to implicitly model the driver's 'reward function' or utility, predicting behaviour based on what action would likely maximise the driver's goals (e.g., balancing progress towards a destination with the costs of braking, potential collision risk, and social considerations like politeness or adherence to norms). The successful performance of this component model suggested that pedestrians may be possessing an intuitive understanding of vehicle dynamics and an intuitive understanding of driver goals and rationality relevant to traffic interactions. This aligns with Theory of Mind (ToM) frameworks where observers infer agents' hidden mental states (goals, beliefs) by inverting a generative model that assumes rational action (Baker et al., 2009; Dindo et al., 2011; Pezzulo et al., 2013; Whiten, 1991). Value-maximisation approaches have been increasingly

applied in human-robot and human-AV interaction research (Kollmitz et al., 2020; Lin et al., 2022; Lio et al., 2020; Markkula et al., 2023; Wang et al., 2015, 2025). These models often incorporate similar factors to the Vb model, such as efficiency, safety, comfort, and social preferences (e.g., inferring selfishness or altruism; Bansal et al., 2020), risk attitudes (Kwon et al., 2020), or proxemic utility (Camara and Fox, 2022). A key distinction, however, is that much of the human-robot interaction literature has been focused on modelling the human from the robot's perspective to inform robot planning (Camara et al., 2021). This thesis offered the complementary inverse perspective by developing and validating a cognitive model of how the pedestrian interprets the vehicle's actions. This perspective has not been commonly modelled but is equally important for designing AVs whose behaviour is interpretable and trustworthy from the human perspective, a point in the section on practical implications further below.

The finding that tau-dot ($\dot{\tau}$) served as a better observational cue within the Ob model than raw deceleration rate supports research highlighting the importance of optical variables in collision perception, time-to-arrival judgments, and understanding motion (Tian et al., 2023; Wang et al., 2025). Some research has emphasised that visual looming cues, such as the rate of change of the vehicle's optical size ($\dot{\theta}$), as critical inputs affecting gap acceptance, potentially via inducing a sense of collision threat (Tian, 2023). The model success based on the visual cue $\dot{\tau}$, in this thesis, suggested that pedestrians are sensitive not only to the presence or magnitude of the deceleration, but more importantly to the rate of change of the current time-to-arrival. At the same time, this implied an adaptation to the dynamics of the braking manoeuvre. Such sensitivity explains why different deceleration profiles (e.g., constant vs. mixed/two-stage deceleration vs. short slowing) resulted in distinct belief patterns in the experiments (and model predictions), even when leading to the same outcome of the vehicle stopping for the pedestrian. Early and consistent deceleration, which would generate a more stable and informative $\dot{\tau}$ signal, appeared to communicate the vehicle's stopping intent more effectively than later and abrupt braking, aligning with findings on the importance of deceleration timing and consistency (Ackermann et al., 2019) and supported the practical importance of clear kinematic cues. The augmented model's success demonstrated how this plausible perceptual cue could be

integrated with higher-level expectations over time to form and update beliefs about vehicle behaviour.

A significant contribution of this research lied in its detailed, time-resolved investigation of the change of pedestrian beliefs regarding vehicle intentions (e.g., the perceived probability of the vehicle yielding). The approach of the current work differed from the single static outcome measures, such as the final decision to cross the road, reaction-detection times, or ratings of willingness to cross the road, often reported in previous studies (Ackermann et al., 2019; Dey et al., 2019; Giles et al., 2019; Markkula et al., 2018; Pekkanen et al., 2022; Tian et al., 2023). The evolving beliefs measured here were able to capture the dynamic nature of belief formation and updating throughout the duration of a vehicle's approach. The findings indicated that pedestrians do not form instantaneous beliefs; rather, their beliefs are formed and updated over time through a continuous process of behaviour evidence integration. The evolution of beliefs over time, that was observed, aligned with the main idea of the Evidence Accumulation Models (EAMs; Giles et al., 2019; Markkula et al., 2018; Pekkanen et al., 2022; Zgonnikov et al., 2024). EAMs formulate binary decision-making – similar to deciding whether an approaching vehicle will stop or not – as a process during which noisy sensory evidence is accumulated over time towards one of the two possible outcome thresholds, e.g., stopping vs not stopping (Myers et al., 2022; Ratcliff et al., 2016). The evidence driving this accumulation process was derived from the perceived kinematic cues discussed previously (speed, deceleration, distance, etc.). The pedestrians' belief curves could possibly be interpreted as an indication of the drift rate of such an underlying evidence accumulator. For example, as more evidence leaning towards the belief of the stopping behaviour is obtained, the belief probability (P_s) increases, similar to an accumulator drifting towards the stopping boundary (Ma et al., 2025; Myers et al., 2022; Ratcliff et al., 2016). The current research could thus provide empirical grounding for applying EAMs to pedestrian-AV interactions, potentially informing the specific inputs (integrated kinematic cues) and the dynamic output (belief trajectory) required for such models. The novelty of using segmented scenarios to measure and model the continuous change of the pedestrian beliefs lied in capturing this belief evolution within the specific, complex context of pedestrian-vehicle interactions. The output of this integration process could be described as

an evolving internal state – a belief, a level of accumulated evidence, or subjective probability – regarding the vehicle's behaviour. Likely, this evolving internal state is what drives the decision-making process, triggering a road crossing action when the belief in the stopping behaviour surpasses a certain threshold, rather than a single outcome measure. Therefore, studying the temporal dynamics of pedestrian beliefs offered a new perspective into the cognitive mechanisms that possibly translate dynamic perception into observable action, subsequently providing a better understanding of the overall road crossing decision process.

The theoretical framework and empirical findings presented in this thesis aligned significantly with the need for, and demonstration of, large-scale integration of computational psychological theories to explain complex human road user interactions, as advocated by (Markkula et al., 2023). In their work it was argued that understanding behaviours in realistic traffic scenarios demands the development and implementation of integrated models that combine elements such as Bayesian perception, Theory of Mind, and value-based decision-making, which moves beyond using beyond isolated cognitive mechanisms. This thesis supported such approach by combining models of two different behaviour estimation mechanisms. The successful application and adaptation of Bayesian observer models from cognitive science (Baker et al., 2009), and the adoption of components from the computational framework developed by Markkula et al. (2023) to specifically model pedestrian beliefs about vehicle stopping intention represented another novel contribution. In particular, the proposed Ob+Vb model implemented in the current research was built directly upon the principles of “Short-term payoff values”, “Behaviour probabilities given actions”, “Behaviour evidence from estimated behaviour value given actions” and “Behaviour evidence from observation of the other agent” presented in Markkula et al. (2023), and in conjunction with an adaptation of Model 1 from Baker et al. (2009). The specific formulation in this thesis integrated the abovementioned components, to predict pedestrians’ beliefs about an approaching vehicle’s behaviour. A key contribution of this research is the validation of the suggested model against empirical pedestrian belief data from two novel experiments, which were specifically designed to capture the temporal dynamics of pedestrian beliefs across a range of vehicle manoeuvres. Therefore, this thesis not only supported

the need for combined theory-building in psychology using real-world situations like pedestrian-vehicle interaction but also offered a validated task-specific model. To summarise, the proposed model's contribution could be explained by the efficacy it demonstrated when integrating behaviour estimation through observations and reasoned expectations to capture pedestrian beliefs.

While general mechanistic models are valuable, pedestrian behaviour exhibits significant variability both between individuals and across different situations. The theoretical framework employed in this research allowed for exploring several key factors contributing to this variability: inherent prior beliefs, the mapping of behaviour evidence onto belief probability, and the adaptive use of information processing strategies. These topics are visited in the following subsections.

5.3.2 From behaviour evidence to belief probabilities

In addition to the integrated cognitive framework, this thesis also introduced a methodological refinement in the form of an alternative softmax transformation function to map accumulated evidence to belief probabilities. Standard softmax functions have been commonly used in computational models to represent choice probabilities (Markkula et al., 2023; Wright and Leyton-Brown, 2017). However, early model testing in this research indicated that a standard softmax did not fully capture the nuanced relationship between the model's evidence calculations and the empirically observed subjective belief ratings from participants. To address this, a modified softmax function was developed, incorporating parameters B (evidence scaling factor, akin to an inverse temperature term) and M (offset constant or bias, specifically towards the stopping behaviour). The inclusion of these parameters, inspired by Richards' family of growth models and concepts from machine learning, provided greater flexibility. This adapted formulation proved important in achieving a better alignment between the model's predicted beliefs and the actual patterns of human subjective ratings, representing a novel methodological contribution to the modelling of ranked belief states in dynamic interaction scenarios.

The consistent finding that the model parameter M often differed significantly from zero (i.e., a 50/50 probability of stopping or not stopping behaviour) provided support for the existence of prior beliefs or expectations in pedestrians (Vilares and Kording, 2011). Pedestrians do not approach interactions free of biases; instead, they bring prior expectations and knowledge regarding the likely

behaviour of approaching vehicles, particularly about stopping behaviour. This aligned with Bayesian models of cognition, where prior probabilities representing existing knowledge or beliefs are combined with incoming sensory evidence to form updated posterior beliefs (Baker et al., 2005, 2009; Dindo et al., 2011).

The formation of these prior beliefs is likely a process influenced by a variety of factors. Through countless interactions with conventional traffic, pedestrians learn typical driver behaviours associated with specific contexts (e.g., deceleration rate, speed and stopping distance at marked vs. unmarked crosswalks, behaviour near schools; Bella and Silvestri, 2015; Fuller, 1984; Sucha et al., 2017). This learning could also be extended to novel systems, where repeated exposure to AVs or specific eHMIs would shape future expectations (Lee et al., 2022). Additionally, formal knowledge, such as traffic laws dictating right-of-way, contributes to priors (e.g., expecting vehicles to yield at a zebra crossing; Habibovic et al., 2018; Sucha et al., 2017), while explicit instructions, like pre-briefings on how an AV or its eHMI functions, could directly shape priors for subsequent interactions (Liu and Hirayama, 2025). Furthermore, a pedestrian's overall trust in technology, and AVs specifically, acts as a powerful prior (Jayaraman et al., 2019). This trust is malleable, influenced by media portrayals, perceived system reliability and competence, and assessments of safety and risk, while novelty and unpredictability associated with AVs can initially lower trust and create more cautious priors (Rezwana et al., 2025). Cultural norms also play a role, as social and cultural expectations regarding road user etiquette and assertiveness can significantly shape prior beliefs about vehicle stopping behaviour beliefs. So, while the formation of these priors is a varied process, the results of this thesis provided support into how they were shaped in different contexts. An indication of that was the difference in the best fitted value of M between the two experiments. In the second experiment, which included a higher proportion of non-stopping segments compared to the first experiment, the best fitted value for M was lower. This suggested that participants in the second experiment did not have the same prior bias as the participants of the first experiment, but their expectations were based on the specific situation they were in, essentially having a reduced prior bias towards the stopping behaviour due to their current experience.

Moreover, studies comparing pedestrian behaviour across nationalities suggest cultural differences in willingness to cross and potentially different baseline expectations of driver behaviour (Feng et al., 2024). Individual differences further modulate prior beliefs, including demographic factors (age, gender), personality traits (e.g., extraversion, conscientiousness, risk aversion), and individual risk perception modulate prior beliefs (Li et al., 2025). For instance, younger individuals or those with higher education levels often report higher expectations and potentially more optimistic priors regarding AVs (Rezwana et al., 2025). Finally, priors are not static but are likely modulated by the immediate situation, such as the type of crossing (signalized vs unsignalized) (Jayaraman et al., 2019), traffic density, or the behaviour of surrounding pedestrians (social influence; Tump et al., 2020). The implication of these priors is significant: they bias the interpretation of the perceived kinematic cues, especially when those cues are ambiguous or under noisy conditions. A strong prior belief (e.g., “cars rarely stop here”) may require substantial and clear kinematic evidence to the contrary (e.g., significant deceleration) before the pedestrian's belief shifts sufficiently towards the vehicle’s stopping behaviour. The per-participant analysis in Chapter 4 showed that these priors were not uniform across individuals. Specifically, the results indicated a tendency for older participants to have a lower prior bias towards the vehicle’s stopping behaviour. This aligned with the abovementioned literature suggesting that factors like age and experience affect expectations in traffic. These findings demonstrated that M was not only an adjustment to improve the model’s accuracy but also a quantifiable reflection of these prior beliefs.

Beyond prior beliefs, the research identifies another source of variability captured by the model parameter B. The consistent estimation of an average value for B across individuals or conditions suggests a stable mapping function that translates behaviour evidence into the belief probability. Parameter B essentially quantifies the sensitivity of the output (belief probability) to changes in the underlying subjective rating. This finding is possibly connected to the cognitive concept of metacognition, which refers to the human capacity to monitor, evaluate, and regulate one's own cognitive processes and states, including the feeling of confidence or certainty associated with a judgment or decision (Lee and Hare, 2023). Parameter B can be seen as reflecting an aspect

of this metacognitive process: how an individual's internal assessment of their certainty about the vehicle's stopping behaviour is transformed into an externalisable probability or behavioural propensity. Stronger, clearer, and more consistent behaviour evidence supporting one interpretation (e.g., stopping behaviour) over the alternative should lead to a higher degree of internal certainty (Schooler et al., 2024). Conversely, ambiguous cues, conflicting information, or high perceptual noise (e.g., at greater distances) would likely result in lower subjective certainty (Scheller et al., 2025). Importantly, individuals differ in how they report their internal certainty levels (Lee and Hare, 2023). Some individuals might exhibit a conservative translation (high B value), requiring a very high level of internal certainty before reporting a high certainty of believing in a vehicle behaviour. Others might have a more liberal translation (low B value), expressing their belief with high certainty even with moderate or even little available behaviour evidence. These individual differences in the translation from behaviour evidence to belief probabilities could stem from personality factors (e.g., general cautiousness, optimism), cognitive style, or learned response biases; Pallier et al., 2002; Šrol and De Neys, 2021). The theoretical implication was that the reported subjective ratings of vehicle behaviour provided by the participants were not a direct reflection of the behaviour evidence but needed to be transformed. Two pedestrians might perceive the same cues and obtain the same behaviour evidence yet report different subjective ratings due to individual differences. This highlighted an important source of variability that seemed to be separate of prior beliefs. The current research did not distinguish whether this mapping was only in play when providing verbal reports about beliefs or if it is in play also when deciding on crossing. For the purposes of the current work, however, modelling this mapping proved important to accurately predict pedestrians' reported beliefs.

5.3.3 Adaptive cue processing

The findings hinting at a potential shift in the cues pedestrians prioritise, when they update their beliefs regarding the approaching vehicle's behaviour, based on the spatiotemporal distance to the vehicle suggested an adaptive information processing strategy, similar to the one proposed by Tian (2023) regarding road crossing decisions. Specifically, the observation that pedestrians might rely more on heuristics or prior expectations at larger distances, while potentially integrating

more detailed kinematic information like deceleration profiles when the vehicle is closer, points towards dynamic adjustments in how information is utilised (DeLucia, 2015; Tian, 2023; Wang et al., 2025). This observation is tied with established principles about optimal cue integration and dynamic cue weighting in perceptual science (Scheller et al., 2025). Cognitive systems are known to dynamically adjust the influence assigned to different sensory cues based on their perceived reliability or certainty in a given context (Fetsch et al., 2009). In the context of road crossing, visual estimates of vehicle kinematics, particularly acceleration or deceleration, are inherently noisier and less reliable at greater distances (Schmidt et al., 2019; Sripada et al., 2021). A computationally rational system, therefore, should be less influenced by these less reliable cues under such conditions (Fetsch et al., 2009). This might lead pedestrians to rely more heavily on simpler, potentially more stable cues like perceived distance or average speed, or to fall back on pre-existing prior beliefs or learned heuristics (DeLucia, 2015; Tian, 2023). As the vehicle approaches the pedestrian, the reliability of detailed kinematic cues (like deceleration pattern and rate) increases, justifying an increase in their influence in the estimation of the vehicle's behaviour and overall decision process (Wang et al., 2025). The potential reliance on heuristics at larger distances is also noteworthy. Heuristics are cognitive shortcuts or rules-of-thumb that allow for faster, less effortful decision-making, especially under conditions of uncertainty or time pressure (Moussaïd et al., 2011). Examples relevant to pedestrian belief updating might include using a simple distance threshold (e.g., "believe that the vehicle is not stopping if it is within X meters and travels with Y speed") or relying on a default prior assumption (e.g., "vehicles travelling at highway speed rarely stop unexpectedly"). Employing such heuristics when detailed kinematic information is unreliable (i.e., at distance) could be an efficient cognitive strategy (DeLucia, 2015; Tian, 2023; Wang et al., 2025). The formulation and performance of Ob+Vb implemented in this thesis provided a computational representation of this adaptive process. In situations where the deceleration-related cues were unreliable, such as at larger distances or during constant speed approaches where the observation evidence were weak Ob underperformed. In these cases, Ob+Vb was driven mainly by the Vb component, which was the source of expectations or heuristics about rational driver behaviour. This was the reason why Vb and Ob+Vb captured the pedestrian belief patterns where Ob failed. Conversely, in cases where the

vehicle was close to the pedestrian and initiating a clear deceleration, the observation-based evidence for the stopping behaviour were strong and less ambiguous. On the other hand, Vb's evidence regarding the two possible behaviours were not discriminable. In these cases, Ob and Ob+Vb performed better than Vb. The better performance of Ob+Vb, which combined the two behaviour estimation component models, indicated that pedestrians do not rely on a single strategy, but adapt their information processing to the quality and nature of the available cues.

Furthermore, the concepts of prior beliefs, subjective certainty, and dynamic cue weighting are likely interrelated. Prior beliefs (M) establish the baseline expectation. Its influence was strongest at the beginning of a vehicle approach, where it affected the initial belief predictions of Vb before stronger observation-based evidence was accumulated. As the vehicle approach unfolded, the reliability of the perceived cues determined how much the belief was updated away from the prior. Lower cue reliability or in other words behaviour evidence conflict ($A_s - A_{ns} \cong 0$) leads to lower overall certainty. Consequently, a greater relative influence of the prior belief on the reported belief, mediated by the individual's specific evidence-to-belief mapping (B) (Lee and Coricelli, 2020). The observations presented in this subsection could provide a plausible cognitive mechanism underpinning the observed shift towards reliance on priors or heuristics when estimating the behaviour of more distant vehicles.

5.4 Integrating behaviour estimation within the road crossing task and cognitive frameworks

The behaviour estimation mechanisms modelled in this thesis, particularly the augmented model (Ob+Vb), represented a plausible precursor to the broader pedestrian road crossing task. In cognitive science it has been suggested that internal beliefs and intentions guide overt behaviour (Ajzen, 1991). In the current thesis, this principle was investigated in the context of pedestrian-vehicle interactions. For instance, Pekkanen et al. (2022) provided model-based evidence that pedestrians engage in a process akin to intent recognition showing that the inclusion of deceleration-related cues ($\dot{\tau}$) was necessary to explain the

timing of crossing decisions. The current research built upon this, providing a detailed model that explained how this intent recognition, or in this thesis' terms behaviour estimation, occurs through the integration of deceleration-related cues perception (Ob) and rational expectations of the other agent's behaviour (Vb). This estimation of vehicle behaviour was assumed to be an input for subsequent phases of the road crossing task, informing the decision making, action selection and eventually execution, as indicated by the computational framework of Markkula et al. (2023).

Behaviour estimation could also be mapped within established cognitive frameworks, such as Situation Awareness (SA; Endsley, 1995) and Human Information Processing (HIP; Lee et al., 2017). The initial sensory processing stages of the models, where pedestrians receive perceptual cues like $\dot{\tau}$, speed, and distance, correspond to Level 1 SA and the Perceptual Stage of HIP. The computations of Ob+Vb, where these cues are integrated resulting into behaviour evidence which are translated into probabilities regarding the current behaviour of the vehicle, represent Level 2 SA (Comprehension) and the Cognitive Stage of HIP (which includes working memory, decision and response selection). These beliefs, P_S, represent the projection of the vehicle's future state answering to the question "Is the vehicle stopping or not?", aligning with Level 3 SA, and is possibly the precursor to the Action Stage in the HIP model, as described in the paragraph above. So, a confident and early belief that the vehicle is stopping, for example, would potentially lead to a smoother, less pressured crossing. Conversely, uncertainty or a late-forming belief regarding the vehicle's behaviour might lead to hesitation or no action. As stated before, a main contribution of this thesis was providing a validated model (Ob+Vb) that explained how these beliefs are formed and updated, and in this section its cognitive mechanisms' plausibility was briefly discussed.

5.5 Potential neurophysiological underpinning of behaviour estimation

The behaviour estimation mechanisms proposed in this thesis, particularly the augmented model (Ob+Vb), exhibited significant similarities and potential

neurophysiological underpinnings in the operation of the Mirror Neuron System (MNS). The MNS is consisted of neurons that discharge when an individual performs a specific action and but importantly also when they observe another individual performing a similar action (Jeon and Lee, 2018; Kilner et al., 2007a, 2007b; Oberman et al., 2007; Proverbio and Zani, 2023; Rizzolatti and Craighero, 2004). This system was first discovered in the premotor cortex (F5) of macaque monkeys (di Pellegrino et al., 1992; Rizzolatti et al., 2001), while Molenberghs et al. (2012) and Rizzolatti and Sinigaglia (2016) findings suggested that MNS properties are also present in the human brain. This unique characteristic has led researchers to posit the MNS as a fundamental neural underpinning for understanding the actions and intentions of others (Zhao et al., 2024), an outcome equivalent to what the behaviour estimation mechanisms, investigated in this thesis, have been described to do.

The ability of humans to infer the intentions of others through the observation of their actions is basic for social cognition and the MNS is considered a plausible candidate for mediating this 'mind-reading' ability (Kilner et al., 2007b). However, the precise mechanism by which intentions are inferred from observed movements is complex, especially given that similar movements can arise from different underlying goals or intentions (Kilner et al., 2007a). This is analogous to the challenges faced by pedestrians in interpreting an approaching vehicle's behaviour, where similar kinematic cues might precede either a stopping or non-stopping behaviour.

Kilner et al. (2007b) proposed that the MNS solves this ill-posed problem through a predictive coding framework based on empirical Bayesian inference. Within this framework, the most likely cause of an observed movement (i.e., the underlying intention or goal) is inferred by minimising prediction error across all relevant cortical levels. This is similar to how Ob+Vb works, where pedestrian beliefs about vehicle behaviour are updated by integrating direct observations (akin to sensory input in the MNS) with value-based expectations (akin to prior predictions or generative models of rational driver behaviour). The MNS, therefore, is not simply a passive reflection of observed actions but an active inferential system that predicts the sensory consequences of an observed agent's motor commands based on an expectation of their goal and then uses prediction error to update these inferences. Therefore, the main finding of this thesis, that

an integrated model (Ob+Vb) is required to accurately predict pedestrian beliefs, provides computational and behavioural support to the idea that action understanding in this context is an active, inferential process, rather than a passive representation of vehicle motion.

Furthermore, the MNS is not solely driven by visual input; it also activates when the sight of a movement is partly occluded, suggesting it predicts the most likely kinematics regardless of full visibility (Kilner et al., 2007b). Research has shown that the human MNS is sensitive not only to the physical aspects of an action but also to the underlying intentions and the social context (Oberman et al., 2007). For instance, activity in the inferior frontal gyrus (IFG), a key area of the human MNS, is modulated by the inferred purpose of an observed action (Oberman et al., 2007). The MNS' involvement in social cognition is further highlighted by its proposed role in empathy and theory of mind. Studies using Electroencephalography (EEG) mu wave suppression as an index of MNS activity (Fox et al., 2016; Proverbio and Zani, 2023) have shown that mu suppression is controlled by the degree of social interaction observed, with more interactive, and socially and contextually relevant stimuli resulting in greater MNS activity (Oberman et al., 2007; Proverbio and Zani, 2023). This suggests a specialisation of the human MNS for processing socially relevant stimuli (other agents' actions), which is relevant to pedestrian-vehicle interactions as pedestrians are essentially interpreting the social affordances (Orban et al., 2021) offered by the approaching vehicle's behaviour.

In essence, just as the MNS allows an observer to transform visual information about another's actions into knowledge about their internal state (e.g., intentions), the behaviour estimation mechanisms modelled in this thesis allow a pedestrian to transform kinematic information from an approaching vehicle into a belief about its future behaviour. The proposed Ob+Vb model, therefore, can be viewed as a computational-level description (Marr, 1982) of the processes that might be implemented at the neuronal level by the MNS and associated brain regions during the pedestrians' estimation of the approaching vehicle's behaviour. Moreover, as advocated by Marr (1982) and echoed more recently by Niv (2021), an important objective of computational cognitive science is to successfully develop computational models that explain behaviour. The validation of such a behavioural model does not require a direct mapping to its neurophysiological

underpinning, even though achieving that mapping would be an extra step towards gaining a better understanding of the cognitive mechanisms involved. Thus, the aim of the work presented in this thesis was primarily to establish a computationally sound and behaviourally validated model of behaviour estimation, rather than to investigate the underlying neurobiological structure per se.

5.6 Practical implications

The findings and the developed Ob+Vb model have several practical implications for the design and deployment of automated vehicles, the development of driver assistance systems and the enhancement of road safety.

5.6.1 Designing interpretable AV behaviours

A key practical implication of these findings is the need to design AVs to generate behaviours that are not just safe but also easily and accurately interpretable by pedestrians. Simply avoiding collisions is insufficient; AV behaviour must align with human cognitive processes and expectations to foster trust and efficient interactions (Rezwana and Lownes, 2024). The better performance of Ob+Vb emphasised that pedestrians interpret vehicle actions through a combination of observed motion and assumptions about rational behaviour. AVs should therefore be designed with this combined mechanism in mind. Based on the findings of this thesis, the following recommendations for AV kinematic behaviour can be made.

1) AVs should utilise clear, consistent, and timely kinematic cues. The experiments demonstrated that early and noticeable deceleration is more effective in signalling stopping intent than late, harsh braking, even if the latter involves higher peak deceleration rates. This aligns with findings suggesting defensive deceleration profiles are preferred and lead to earlier crossing initiation (Ackermann et al., 2019; Dietrich et al., 2020). Constant speed approach phases should be unambiguous, and unpredictable or overly subtle changes in speed should be avoided, especially when nearing pedestrians, as these create uncertainty that hinders accurate

belief formation. Designing AV motion planners to explicitly generate such interpretable kinematic profiles is therefore important (Moller et al., 2025).

2) AVs require sophisticated algorithms to predict pedestrian behaviour for safe planning (Camara et al., 2020b). This thesis provides the inverse perspective: a model of how pedestrians predict vehicle behaviour. This understanding is also important for AVs. If an AV can use a model analogous to Ob+Vb to simulate a pedestrian's beliefs in real-time, it could predict whether its planned manoeuvre is likely to be perceived as clear, ambiguous, or even misleading, and adapt its plan accordingly (Prédhumeau et al., 2022). For instance, if the model predicts that a planned deceleration is too subtle to be interpreted as a clear intention to yield, the AV could proactively adjust its trajectory to be clearer, perhaps by decelerating slightly earlier or more distinctly. This would possibly allow for more proactive and cooperative interactions.
3) The partial success of the value-based component (Vb) suggests that designing AVs to behave in ways that align with pedestrians' expectations of rational, goal-directed, and socially considerate behaviour (e.g., showing appropriate politeness or caution based on context) could enhance interpretability and trust (Camara and Fox, 2022). This involves programming AVs with behaviours that reflect an understanding of implicit road rules and are in accordance with the appropriate social norms.
4) A further consideration arising from these implications involves a potential long-term feedback loop. If cognitive models like Ob+Vb are successfully used to design AVs that exhibit highly interpretable and predictable behaviour, the task of interpreting these AVs may become cognitively simpler for pedestrians over time. As pedestrians learn the consistent behavioural patterns of AVs, their reliance on complex inferential processes might decrease. In such a future, simpler cognitive models – perhaps closer to the Ob or Vb components alone, with strong, learned priors reflecting established AV norms – might become sufficient to accurately predict AV intentions. This suggests a potential co-evolution where human-centred AV design, informed by current cognitive models, could eventually lead to reduced cognitive load for pedestrians and potentially shift the dominant cognitive strategies employed in these interactions.

5.6.2 Informing driver assistance systems and eHMI design

While the primary focus is on AVs, these principles are also relevant for advanced driver assistance systems (ADAS) in human-driven vehicles. ADAS that influence vehicle kinematics should do so in a manner that is clearly interpretable to nearby pedestrians.

1) For systems like Adaptive Cruise Control (ACC) that may react to pedestrians, the findings on effective kinematic cues (e.g., early, consistent deceleration) are directly applicable. An ACC that decelerates smoothly and noticeably for a pedestrian far in distance communicates its intent better than braking later and more sharply.
2) Furthermore, the Ob+Vb model could be used to evaluate the legibility of ADAS-controlled vehicle behaviours for pedestrians, ensuring that automated assistance enhances rather than complicates pedestrian-vehicle interactions.
3) Although this thesis focused on implicit communication via vehicle kinematics, the findings could inform eHMI design. The proposed model highlights scenarios where pedestrian belief uncertainty is high, typically when implicit kinematic cues are weak or ambiguous (e.g., a distant vehicle or one moving at a slow, constant speed). These are scenarios where an explicit signal from an eHMI could be most beneficial to resolve ambiguity and clarify intent (Lau et al., 2022). The Ob+Vb model could potentially be used to identify these ambiguous cases in real-time and trigger an appropriate eHMI display. While pedestrians may not feel at risk from a distant vehicle, an early explicit signal could still be valuable for establishing trust and better comprehension of the AV's plan. However, the literature cautions that eHMI signals should not be I conflict with the vehicle's kinematics to be effective and avoid undermining trust (Lau et al., 2022).

5.6.3 Enhancing simulation and safety

Beyond informing the design of vehicle behaviours, the computational framework developed in this thesis has practical applications as a tool for improving how AV systems are tested and how road safety is analysed.

1) The development and testing of AV systems rely on simulation environments (Zhao et al., 2024). Ob+Vb could contribute to creating more

realistic virtual pedestrians (agents) for these simulations. Current simulators like CARLA (CARLA, 2025) or SUMO (Rampf et al., 2023), and co-simulations thereof (Gutiérrez-Moreno et al., 2022), often employ descriptive or rule-based pedestrian models which are not offering the interpretability and more in-depth explanations of the cognitive processes underlying behaviour. Implementing agents whose belief states evolve according to cognitive principles like those captured by Ob+Vb would allow these agents to react to simulated AV behaviour in a more human-like manner, providing a more valid and challenging testbed for AV algorithms compared to agents reacting based on rules or predefined settings (Wang et al., 2023). While implementing complex cognitive models in real-time, large-scale simulations presents challenges (Moller et al., 2025), the principles derived from Ob+Vb can guide the development of higher-fidelity cognitive agents.

2) Understanding the root causes of pedestrian-vehicle incidents often involves understanding misinterpretations of intent (Alambeigi et al., 2020; Habibovic and Davidsson, 2012). The Ob+Vb model could be used retrospectively to analyse accident scenarios or near-misses, identifying specific kinematic patterns or interaction sequences that are prone to misinterpretation by pedestrians due to conflicting perceived kinematic cues and rational expectations. Ob+Vb provides a more psychologically grounded basis for how pedestrian agents in microscopic traffic simulation models estimate vehicle behaviour, thus it could offer more accurate predictions of overall pedestrian behaviour, conflict likelihood and the effectiveness of safety interventions within these simulations. This thesis, by detailing how beliefs are formed and updated, could contribute to new metrics for assessing the interpretability/predictability/effectiveness of vehicle manoeuvres or eHMI designs from the pedestrian's cognitive viewpoint. Eventually, these metrics could become part of safety assessment protocols for new vehicle systems or infrastructure changes. Another implication could involve infrastructure design (e.g., crosswalk or traffic light placement) or targeted driver/pedestrian education programs.

5.6.4 Experimental framework

Another significant practical implication of the current thesis is its contribution to experimental methodology. The work details and validates a robust experimental framework designed to investigate the behaviour estimation mechanisms of pedestrians while they observe approaching vehicles. This framework, validated across two extensive experiments, used an immersive virtual reality environment to present pedestrians with a variety of controlled yet realistic vehicle approach scenarios, including the systematic manipulation of vehicle speed, TTA, and diverse driving manoeuvres. The experimental design, adapted from cognitive science paradigms (Baker et al., 2009), successfully captured the temporal evolution of pedestrian beliefs by presenting scenarios in truncated segments. At each judgment point, a two-part response (behaviour judgment and confidence rating) was employed to quantify the belief state and its associated certainty. The success of this framework in generating nuanced data capable of stringently testing computational models like the Ob+Vb, and its proven utility in assessing model generalisability, establishes it as a robust and replicable paradigm for future research into pedestrian perception, cognition, and interaction with both human-driven and automated vehicles.

5.7 Limitations

The work provided in this thesis, while providing valuable insights into pedestrian behaviour estimation, is subject to limitations that should be considered when interpreting the findings and planning future work. These limitations can be broadly categorized by their origin: (1) the experimental design, (2) the model assumptions, and (3) the focus on only a cognitive subpart of the whole road crossing task.

The experiments were conducted in an immersive VR environment. While this offers significant advantages in terms of experimental control, safety, and the ability to systematically manipulate complex scenarios, the specific experimental design involving scenarios truncated into segments for belief rating cannot be directly replicated in real-world field studies. Although the VR-based approach allows for true interaction and is superior to simpler methods like video stimuli,

potential limitations remain, including possible distortions in perceptual fidelity (e.g., estimation of speed and distance) and an attenuated sense of risk, which might lead to behavioural differences compared to real-world actions (Wynne et al., 2019). The direct transferability of exact parameter values from VR to real-world behaviour thus requires careful consideration, though the structural findings regarding cognitive mechanisms are likely more robust.

Furthermore, the tested scenarios represent a limited subset of real-world encounters. The experiments focused on interactions with a single approaching vehicle on a straight single-lane road under clear daylight conditions, without the complexities of multiple road users, varied environmental conditions (weather, lighting), or diverse road layouts (curves, complex intersections). Notably, the scenarios did not include explicit communication cues from the vehicle, such as flashing headlights or eHMIs. Moreover, there was no visible driver figure (even virtual) in the approaching vehicle, which might influence pedestrian expectations and trust. The impact of these explicit communication cues and the social presence of a driver on behaviour estimation remains to be explored within this framework.

The participant sample for each experiment consisted of 30 adults. While covering a range of ages, participants under 18 years old were not included, due to ethical considerations. Moreover, participants over 50 years old constituted a small minority of the sample, and in this older subgroup, all were men, limiting representation of and generalisability to older female pedestrians. Some might argue that 30 participants per experiment is a relatively small sample for generalising cognitive models broadly. However, the studies based on which the current experiments were developed (Ackermann et al., 2019; Baker et al., 2009; Dey et al., 2019; Petzoldt et al., 2018), had a similar number of participants. Another limitation relates to the data collection strategy for each participant. By collecting only one belief rating, per segment, the design prioritised testing a wide range of scenarios over assessing the consistency of the judgments for each participant.

A second category of limitations relates to the simplifying assumptions made by the computational models. The Ob component, for instance, relied on $\dot{\tau}$ as the primary perceptual cue for deceleration. Pedestrians possibly integrate a richer set of visual information not captured by the model, potentially including subtle

changes in the vehicle's pitch during braking or minor trajectory deviations. Similarly, Vb assumes a specific reward structure for the driver (e.g., balancing progress, comfort, politeness). Real driver motivations are likely more complex and context dependent. The combination weights of the two evidence sources (β_O, β_V) in the Ob+Vb model were treated as static parameters. The relative weight of observation versus value-based evidence might change based on factors like cue reliability, distance, uncertainty, or cognitive load, a complexity not fully realised in the current model.

The composite belief score, which transformed the binary choice and the confidence rating into a single continuous variable, is susceptible to confidence-based biases and could be considered a methodological limitation. This approach risks misrepresenting the central tendency of the group's opinion, particularly in ambiguous scenarios where the population's belief is near the 50% uncertainty threshold. For instance, a small number of overconfident individual outliers (i.e., those assigning high confidence to a minority choice) could lead to a shift on the aggregated average, effectively masking the agreement between the less confident majority. While this effect may not entirely invalidate the model's overall predicted rank correlation, it hinders the significance of the output as a true representation of the average group beliefs. A more robust approach would be to decouple the decision (choice) and the confidence (certainty). The model would first capture the binary choices of the participants, which Tian et al. (2023) showed that they follow similar patterns to the composite belief scores in the data of the current thesis. The confidence could then be modelled as a separate metric of uncertainty, mitigating the risk of a single participant skewing the overall aggregated belief. Other ways of reducing that risk would involve increasing the data density, either by expanding the dataset (more participants) or through repeated measures to assess the consistency of individual confidence ratings, and/or outlier treatment.

The current modelling approach focused on capturing the continuous dynamics of pedestrian belief updating over time, as influenced by observed and expected kinematics, rather than investigating the accuracy of the pedestrians' judgments. An alternative approach, would be to assess the judgment accuracy, classifying each trial by whether the reported belief correctly inferred the vehicle's true behaviour. While informative, this approach was beyond the scope of this

thesis. The challenge would be to classify beliefs that sit near the 50% uncertainty threshold. Treating these near-50% beliefs as either 'correct' or 'wrong' would impose an arbitrary binary classification on a continuous cognitive state, potentially misrepresenting the pedestrians' true uncertainty and could possibly hide the behaviour estimation mechanisms that this thesis aimed to model. Thus, the choice was made to model the mechanisms of inference across all states of certainty, rather than focusing on a potentially simplified classification of judgment accuracy.

While this thesis addresses noise in perception (i.e., Gaussian noise in observations for the Ob model), more complex noise considerations, such as value-transformed sensory noise, a term borrowed from Markkula et al. (2023), or specific assumptions about where noise is injected, are not made. The models consider a limited set of vehicle behaviours (stopping/not stopping) and pedestrian responses (belief rating). Although, the Markkula et al. (2023) computational framework outlines a more granular approach to action selection based on motor primitives and accumulated action value estimates, it was deemed to be a level of detail beyond the scope of the current belief estimation models.

In addition to these assumptions, limitations also arise from the fitting and validation of the models. Despite taking measures such as Bootstrap Cross-Validation (BSCV) to assess generalisability and mitigate overfitting (as discussed in Chapters 3 and 4), the complexity of Ob+Vb, with its number of parameters, inherently carries a risk of fitting noise in the specific datasets used. Continuous validation on diverse, unseen datasets is important. While per-participant fitting was explored (Chapter 4), confirming the superiority of Ob+Vb at the individual level, analyses prior to that relied on data averaged across participants. Averaging can potentially mask significant individual variability in how pedestrians interpret vehicle behaviour and may obscure the possibility that different subgroups employ qualitatively different strategies, something which should thus be considered not least in the interpretation of the Chapter 3 findings.

Finally, the third type of limitation has to do with the fact that the current research concentrated on modelling the pedestrian's belief about the vehicle's stopping behaviour (P_s). As also discussed in Section 5.4, in the comprehensive framework of Markkula et al. (2023), behaviour estimation is one component of a

larger perception-action loop. The models of this thesis primarily address the “Bayesian perception” and “Behaviour estimation" components leading to the belief P_s, analogous to $P_{b|a}$, in Markkula et al.'s (2023) terms. While this belief is an important input, the subsequent stages of the road crossing task – such as the explicit decision to cross (often modelled via evidence accumulation to a decision threshold), motor planning, and the physical execution of the crossing action (involving motor primitives) – are not explicitly modelled in the current thesis. Factors like risk perception, urgency, specific gap acceptance thresholds, and game-theoretic adaptations to the other agent's unfolding actions are part of this latter stage. A complete model of road crossing would need to integrate the proposed belief estimation models with these subsequent decision-making and action components.

5.8 Future research directions

This thesis has provided valuable insights into the cognitive mechanisms underlying pedestrian behaviour estimation when interacting with approaching vehicles, resulting in the implementation and validation of the augmented Ob+Vb model. However, the findings and methodologies presented herein also open several promising avenues for future investigation to further refine the overall understanding of behaviour estimation mechanisms in the pedestrian-vehicle interaction context and enhance the practical applicability of this work.

A next step would possibly involve validating the Ob+Vb model and its parameterisation in more naturalistic settings. This could include studies in CAVE-based pedestrian simulators but extracting the driving manoeuvres (to be later truncated into segments) either from high-fidelity driving simulators or real-world driving, that would allow for even more realistic behaviour estimation testing. An alternative method to increase realism would be to capture 3D videos of real traffic encounters to be experienced within the CAVE environment, potentially offering a higher degree of visual and contextual fidelity. Such studies are important for confirming the model's real-world applicability and fine-tuning its parameters. Such studies would help to investigate behaviour estimation in a wider array of complex traffic scenarios. This includes exploring interactions at

intersections, on curved roadways, in the presence of multiple road users (e.g., other vehicles, cyclists), and under diverse environmental conditions such as varying weather and lighting. The influence of infrastructure, like marked crosswalks or traffic signals, also warrants investigation.

Fitting Ob+Vb to video data, in contrast to the kinematic data extracted from the controlled VR environment, would require overcoming the challenge of kinematic extraction from visual input. The first step would be a Computer Vision/Deep Learning model to reliably estimate the target vehicle's key kinematic variables from the raw pixel data. The extracted kinematics would provide the inputs to the component models. Finally, since real-time subjective beliefs cannot be obtained from video data, a method to extract pedestrian beliefs, such as asking third-person observer participants to provide continuous belief ratings would be needed to calculate the target output (correlation and/or RMSE between the model predictions and the extracted beliefs).

Further enhancement and integration of the current models represent another important direction. The current models could be extended to incorporate a richer set of observed cues. This might include vehicle trajectory information, changes in vehicle orientation, or even non-kinematic cues such as turn signals or explicit eHMI displays, if present. An interesting avenue for future work is the integration of the validated Ob+Vb model with established road user interaction models to create comprehensive models of the pedestrian behaviour (Markkula et al., 2023). Such integration could happen with models such as Evidence Accumulation Models (EAMs), where the Ob+Vb's belief output P_s could be used to set the drift rate of an Evidence Accumulation Model, where a stronger belief in the vehicle stopping would accelerate the accumulation of evidence toward making a crossing decision. Furthermore, the Ob+Vb currently models the belief of a vehicle's behaviour from the pedestrian's perspective. To integrate this into a full pedestrian-vehicle framework, two adaptations would be needed. First, both agents' beliefs about each other would have to be modelled. Second, the behaviour estimation output (belief about the other's behaviour) would have to be connected to an evidence accumulation step as described above, to complete the decision loop (for both agents). Alternatively, Ob+Vb could be integrated with the Perceptually Plausible Road Crossing Decision (PT-PRD) model as presented in Chapter 6 of Tian (2023) to yield a more comprehensive pedestrian

crossing decision model. The main idea would be to use the belief output from Ob+Vb (P_s) to inform the PT-PRD model's dynamic decision-making processes. This would enhance the PT-PRD model by making it sensitive also to the expectation-based factors that the current belief model captures, potentially improving predictions of crossing initiation timing.

The validated Ob+Vb model provided a first thorough investigation of behaviour estimation in single-vehicle, single-pedestrian interactions. However, in real-worlds traffic scenarios multiple agents might be involved. Thus, an interesting future direction would be to extend the suggested model in a hierarchical manner so that it could integrate social influences and contextual complexity. The presence of other pedestrians could introduce social influences that affect both the perceived intent of the approaching vehicle and the pedestrian's own urgency. The model would need to account for pro-social cues, such as the phenomenon of group size increasing the likelihood of yielding (Park et al., 2024). This could be modelled within the Vb component by making the politeness constant an adaptive parameter that increases with the observed number of attentive pedestrians waiting to cross, thereby affecting the calculated utility for the stopping behaviour. Furthermore, the model could address collective decision-making, where a pedestrian's prior beliefs might be influenced by the actions of nearby pedestrians, i.e., a form of social influence or herd behaviour (Faria et al., 2010) that could impact the confidence of the initial estimation. Conversely, the presence of other vehicles (e.g. non-conflicting traffic, traffic flow behind the approaching vehicle) would introduce contextual ambiguity and additional constraints for the driver/AV. In the Ob component, the visual presence of surrounding traffic could increase cognitive load and visual clutter, which may be modelled by increasing the perceptual noise and/or lowering the weight assigned to the observation-based evidence. For the Vb component, calculating the utility would also be more complex, since there would be the need to include a cost term for traffic impedance, i.e., the penalty for unnecessary or unsafe deceleration that disrupts the flow of following vehicles.

The potential for dynamic weighting of the observation-based and value-based components within the Ob+Vb model should be explored. The weighting of these mechanisms' evidence might shift based on contextual factors like cue

reliability, distance to the vehicle, perceived ambiguity, or the pedestrian's cognitive state.

Investigating individual and group differences would also be interesting since a uniform model fit would likely be insufficient for capturing all pedestrians' beliefs. Drivers and AVs should interact safely with a diverse population, including older adults or children, whose beliefs of vehicle behaviour may differ from an "average" adult. Therefore, larger-scale studies would be beneficial to systematically explore how factors such as age, gender, risk tendency, cognitive abilities, trust in automation, prior experiences, and cultural background influence the parameters of Ob+Vb (particularly the prior bias M and the belief mapping B) and the overall accuracy of its predictions. Correlating fitted model parameters with validated psychometric measures could provide deeper insights and potentially lead to personalized models or AV interaction strategies tailored to different user groups. For instance, one could investigate if scores on a risk-taking or a trust in automation scales (Blais and Weber, 2006; Kohn et al., 2021; Zhang et al., 2019) correlate with the model's prior bias or belief mapping parameters. This could lead to more personalised models or AV interaction strategies associated to different road user groups.

The practical value of the Ob+Vb model lies in its potential application in AV development, simulation, and broader road safety analysis. As discussed in Section 5.6 future engineering research should focus on implementing these insights into AV perception and planning systems. This includes designing AVs that generate interpretable kinematics and potentially equipping AVs with an inverse model based on Ob+Vb to anticipate pedestrian interpretations, leading to safer and more comfortable interactions. Furthermore, the cognitive principles captured by Ob+Vb can inform the development of more realistic virtual pedestrian agents for AV testing in simulation environments like CARLA or SUMO. Creating smart agents whose beliefs evolve according to these principles would provide a more valid testbed for AV algorithms compared to current rule-based agents.

Future work could also delve into the neurophysiological underpinnings of behaviour estimation. Neurophysiological studies, perhaps employing EEG to measure mu suppression (Fox et al., 2016; Proverbio and Zani, 2023) or fMRI (Caspers et al., 2010; Molenberghs et al., 2012) to identify active regions during

simulated behaviour estimation tasks similar to those used in this thesis, could investigate the neural correlates of the belief updating processes modelled by Ob+Vb. This could provide direct evidence for the engagement of MNS during behaviour estimation in road traffic, offering a more holistic understanding and link the computational mechanisms of behaviour estimation to their neural implementation.

Finally, building on the concept of a long-term feedback loop discussed previously (Section 5.6), another interesting research path is to explore how pedestrian behaviour estimation and trust evolve over long-term exposure to AVs. As a reminder, if AVs become more predictable, the cognitive load on pedestrians might decrease, potentially shifting the dominant cognitive strategies they employ.

By pursuing these future research directions, the understanding of pedestrian cognitive mechanisms can be further advanced, contributing to the development of safer road environments and more intuitive and socially adept interactions between humans and AVs.

5.9 Concluding remarks

The work of this thesis presents a novel comprehension of how pedestrians estimate the behaviour of approaching vehicles to eventually make road crossing decisions. This research represents the first attempt to implement and validate computational models of the specific cognitive mechanisms underlying this process. The results demonstrate that pedestrian beliefs regarding the behaviour of an approaching vehicle are not based solely on perceiving deceleration-related cues or on rational, value-maximising expectations alone, but requires an integrated framework combining both. The implementation and validation of a successful behaviour estimation model has been demonstrated by adopting and combining established cognitive science models, and by adapting cognitive science simplified experimental laboratory paradigms to the pedestrian-vehicle interaction setting. From this thesis' findings, the augmented Ob+Vb model can be suggested as a plausible psychological and successful computational explanation for how pedestrians estimate vehicle behaviours. This integrated

behaviour estimation mechanism could allow for the improvement of human-AV interaction design and the development of more realistic cognitive agents for use in traffic simulation. However, future work should investigate these mechanisms across a wider range of scenarios and populations, to pave the way for the development and deployment of socially capable automated vehicles whose behaviour is intuitively and accurately understood by pedestrians.

References

A. H. Kalantari, Y. Yang, Y. M. Lee, N. Merat, G. Markkula, 2023. Driver-Pedestrian Interactions at Unsignalized Crossings Are Not in Line With the Nash Equilibrium. IEEE Access 11, 110707–110723. https://doi.org/10.1109/ACCESS.2023.3322959

A. R. Srinivasan, Y. -S. Lin, M. Antonello, A. Knittel, M. Hasan, M. Hawasly, J. Redford, S. Ramamoorthy, M. Leonetti, J. Billington, R. Romano, G. Markkula, 2023. Beyond RMSE: Do Machine-Learned Models of Road User Interaction Produce Human-Like Behavior? IEEE Trans. Intell. Transp. Syst. 24, 7166–7177. https://doi.org/10.1109/TITS.2023.3263358

A. Rasouli, I. Kotseruba, J. K. Tsotsos, 2017a. Are They Going to Cross? A Benchmark Dataset and Baseline for Pedestrian Crosswalk Behavior, in: 2017 IEEE International Conference on Computer Vision Workshops (ICCVW). Presented at the 2017 IEEE International Conference on Computer Vision Workshops (ICCVW), pp. 206–213. https://doi.org/10.1109/ICCVW.2017.33

A. Rasouli, I. Kotseruba, J. K. Tsotsos, 2017b. Agreeing to cross: How drivers and pedestrians communicate, in: 2017 IEEE Intelligent Vehicles Symposium (IV). Presented at the 2017 IEEE Intelligent Vehicles Symposium (IV), pp. 264–269. https://doi.org/10.1109/IVS.2017.7995730

A. Rasouli, J. K. Tsotsos, 2020. Autonomous Vehicles That Interact With Pedestrians: A Survey of Theory and Practice. IEEE Trans. Intell. Transp. Syst. 21, 900–918. https://doi.org/10.1109/TITS.2019.2901817

A. T. Schulz, R. Stiefelhagen, 2015a. Pedestrian intention recognition using Latent-dynamic Conditional Random Fields, in: 2015 IEEE Intelligent Vehicles Symposium (IV). Presented at the 2015 IEEE Intelligent Vehicles Symposium (IV), pp. 622–627. https://doi.org/10.1109/IVS.2015.7225754

A. T. Schulz, R. Stiefelhagen, 2015b. A Controlled Interactive Multiple Model Filter for Combined Pedestrian Intention Recognition and Path Prediction, in: 2015 IEEE 18th International Conference on Intelligent Transportation Systems. Presented at the 2015 IEEE 18th International Conference on Intelligent Transportation Systems, pp. 173–178. https://doi.org/10.1109/ITSC.2015.37

Ackermann, C., Beggiato, M., Bluhm, L.-F., Krems, J., 2018. Vehicle movements as implicit communication signal between pedestrians and automated vehicles.

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. Part F Traffic Psychol. Behav. 62, 757–768. https://doi.org/10.1016/j.trf.2019.03.006

Ahmed, S., Huda, M.N., Rajbhandari, S., Saha, C., Elshaw, M., Kanarachos, S., 2019. Pedestrian and Cyclist Detection and Intent Estimation for Autonomous Vehicles: A Survey. Appl. Sci. 9, 2335. https://doi.org/10.3390/app9112335

Ajzen, I., 1991. The theory of planned behavior. Organ. Behav. Hum. Decis. Process., Theories of Cognitive Self-Regulation 50, 179–211. https://doi.org/10.1016/0749-5978(91)90020-T

Alambeigi, H., McDonald, A.D., Tankasala, S.R., 2020. Crash Themes in Automated Vehicles: A Topic Modeling Analysis of the California Department of Motor Vehicles Automated Vehicle Crash Database. ArXiv Appl.
Amoruso, L., Urgesi, C., 2016. Contextual modulation of motor resonance during the observation of everyday actions. NeuroImage 134, 74–84. https://doi.org/10.1016/j.neuroimage.2016.03.060
Antić, B., Pešić, D., Milutinović, N., Maslać, M., 2016. Pedestrian behaviours: Validation of the Serbian version of the pedestrian behaviour scale. Transp. Res. Part F Traffic Psychol. Behav. 41, 170–178. https://doi.org/10.1016/j.trf.2016.02.004
Ashby, F.G., Maddox, W.T., Lee, W.W., 1994. On the Dangers of Averaging across Subjects When Using Multidimensional Scaling or the Similarity-Choice Model. Psychol. Sci. 5, 144–151.
Baker, C., Saxe, R., Tenenbaum, J., 2011. Bayesian Theory of Mind: Modeling Joint Belief-Desire Attribution. Proc. Annu. Meet. Cogn. Sci. Soc. 33.
Baker, C., Saxe, R., Tenenbaum, J., 2005. Bayesian models of human action understanding, in: Advances in Neural Information Processing Systems. MIT Press.
Baker, C.L. (Chris L., 2012. Bayesian Theory of Mind : modeling human reasoning about beliefs, desires, goals, and social relations (Thesis). Massachusetts Institute of Technology.
Baker, C.L., Jara-Ettinger, J., Saxe, R., Tenenbaum, J.B., 2017. Rational quantitative attribution of beliefs, desires and percepts in human mentalizing. Nat. Hum. Behav. 1, 1–10. https://doi.org/10.1038/s41562-017-0064
Baker, C.L., Saxe, R., Tenenbaum, J.B., 2009. Action understanding as inverse planning. Cognition, Reinforcement learning and higher cognition 113, 329–349. https://doi.org/10.1016/j.cognition.2009.07.005
Bansal, S., Xu, J., Howard, A., Isbell, C., 2020. A Bayesian Framework for Nash Equilibrium Inference in Human-Robot Parallel Play. https://doi.org/10.48550/arXiv.2006.05729
Bao, J., n.d. Lectures 8: Two-, Three-, and Four-Way ANOVA.
Bay, Y.Y., Yearick, K.A., 2024. Machine Learning vs Deep Learning: The Generalization Problem. https://doi.org/10.48550/arXiv.2403.01621
Bazilinskyy, P., Dodou, D., de Winter, J., 2019. Survey on eHMI concepts: The effect of text, color, and perspective. Transp. Res. Part F Traffic Psychol. Behav. 67, 175–194. https://doi.org/10.1016/j.trf.2019.10.013
Beggiato, M., Witzlack, C., Krems, J.F., 2017. Gap Acceptance and Time-To-Arrival Estimates as Basis for Informal Communication between Pedestrians and Vehicles, in: Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI '17. Association for Computing Machinery, New York, NY, USA, pp. 50–57. https://doi.org/10.1145/3122986.3122995
Bella, F., Silvestri, M., 2015. Effects of safety measures on driver's speed behavior at pedestrian crossings. Accid. Anal. Prev. 83, 111–124. https://doi.org/10.1016/j.aap.2015.07.016
Bennett, S., Felton, A., Akcelik, R., 2002. Pedestrian movement characteristics at signalised intersections.
Berry, J.W., Schwebel, D.C., 2009. Configural approaches to temperament assessment: Implications for predicting risk of unintentional injury in

children. J. Pers. 77, 1381–1409. https://doi.org/10.1111/j.1467-6494.2009.00586.x
Blais, A.-R., Weber, E.U., 2006. A Domain-Specific Risk-Taking (DOSPERT) scale for adultpopulations. Judgm. Decis. Mak. 1, 33–47. https://doi.org/10.1017/S1930297500000334
Bock, J., Kotte, J., Beemelmanns, T., Klösges, M., 2024. Self-learning Trajectory Prediction with Recurrent Neural Networks at Intelligent Intersections. Presented at the 3rd International Conference on Vehicle Technology and Intelligent Transport Systems, pp. 346–351.
Böckle, M.-P., Brenden, A.P., Klingegård, M., Habibovic, A., Bout, M., 2017. SAV2P: Exploring the Impact of an Interface for Shared Automated Vehicles on Pedestrians' Experience, in: Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications Adjunct, AutomotiveUI '17. Association for Computing Machinery, New York, NY, USA, pp. 136–140. https://doi.org/10.1145/3131726.3131765
Bokare, P.S., Maurya, A.K., 2017. Acceleration-Deceleration Behaviour of Various Vehicle Types. Transp. Res. Procedia, World Conference on Transport Research - WCTR 2016 Shanghai. 10-15 July 2016 25, 4733–4749. https://doi.org/10.1016/j.trpro.2017.05.486
Brewer, M.A., Fitzpatrick, K., Whitacre, J.A., Lord, D., 2006. Exploration of Pedestrian Gap-Acceptance Behavior at Selected Locations. Transp. Res. Rec. 1982, 132–140. https://doi.org/10.1177/0361198106198200117
Brill, S., Kumar Debnath, A., Payre, W., Horan, B., Birrell, S., 2024. Factors influencing the perception of safety for pedestrians and cyclists through interactions with automated vehicles in shared spaces. Transp. Res. Part F Traffic Psychol. Behav. 107, 181–195. https://doi.org/10.1016/j.trf.2024.08.032
Brown, B., Laurier, E., 2017. The Trouble with Autopilots: Assisted and Autonomous Driving on the Social Road, in: Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems, CHI '17. Association for Computing Machinery, New York, NY, USA, pp. 416–429. https://doi.org/10.1145/3025453.3025462
Camara, F., Bellotto, N., Coşar, S., Nathanael, D., Althoff, M., Wu, J., Ruenz, J., Dietrich, A., Fox, C.W., 2020a. Pedestrian Models for Autonomous Driving Part I: Low-Level Models, From Sensing to Tracking. IEEE Trans. Intell. Transp. Syst. 22, 6131–6151.
Camara, F., Bellotto, N., Coşar, S., Weber, F., Nathanael, D., Althoff, M., Wu, J., Ruenz, J., Dietrich, A., Markkula, G., Schieben, A., Tango, F., Merat, N., Fox, C.W., 2020b. Pedestrian Models for Autonomous Driving Part II: High-Level Models of Human Behavior. IEEE Trans. Intell. Transp. Syst. 22, 5453–5472.
Camara, F., Dickinson, P., Fox, C., 2021. Evaluating pedestrian interaction preferences with a game theoretic autonomous vehicle in virtual reality. Transp. Res. Part F Traffic Psychol. Behav. 78, 410–423. https://doi.org/10.1016/j.trf.2021.02.017
Camara, F., Fox, C., 2022. Unfreezing autonomous vehicles with game theory, proxemics, and trust. Front. Comput. Sci. 4. https://doi.org/10.3389/fcomp.2022.969194

Cambridge Dictionary, 2024. Meaning of pedestrian in English [WWW Document]. pedestrian. URL https://dictionary.cambridge.org/dictionary/english/pedestrian
CARLA, 2025. CARLA [WWW Document]. CARLA Simulator. URL http://carla.org//
Carlowitz, S., Madigan, R., Goodridge, C.M., Hilz, J., Marberger, C., Alt, P., Schulz, M., Osswalt, S., Engeln, A., Merat, N., 2024. Balancing Comfort Deceleration Rate and Environmental Impact at Crosswalks and Intersections in Automated Driving. https://doi.org/10.2139/ssrn.5063808
Carlowitz, S., Madigan, R., Lee, Y.M., Tango, F., Merat, N., 2023. Pedestrians' perceptions of automated vehicle movements and light-based eHMIs in real world conditions: A test track study. Transp. Res. Part F Traffic Psychol. Behav. 95, 83–97. https://doi.org/10.1016/j.trf.2023.02.010
Caspers, S., Zilles, K., Laird, A.R., Eickhoff, S.B., 2010. ALE meta-analysis of action observation and imitation in the human brain. NeuroImage 50, 1148–1167. https://doi.org/10.1016/j.neuroimage.2009.12.112
Centre for Connected and Autonomous Vehicles, 2022. Connected & Automated Mobility 2025: Realising the benefits of self-driving vehicles in the UK.
Chandra, S., Rastogi, R., Das, V.R., 2014. Descriptive and parametric analysis of pedestrian gap acceptance in mixed traffic conditions. KSCE J. Civ. Eng. 18, 284–293. https://doi.org/10.1007/s12205-014-0363-z
Chandrapp, A.K., Bhattacharyya, K., Maitra, B., 2016. Estimation of Post-Encroachment Time and Threshold Wait Time for Pedestrians on a Busy Urban Corridor in a Heterogeneous Traffic Environment: An Experience in Kolkata. Asian Transp. Stud. 4, 421–429. https://doi.org/10.11175/eastsats.4.421
Chang, C.-M., Toda, K., Sakamoto, D., Igarashi, T., 2017. Eyes on a Car: an Interface Design for Communication between an Autonomous Car and a Pedestrian, in: Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI '17. Association for Computing Machinery, New York, NY, USA, pp. 65–73. https://doi.org/10.1145/3122986.3122989
Chemero, A., 2003. An Outline of a Theory of Affordances. Ecol. Psychol. 15, 181–195. https://doi.org/10.1207/S15326969ECO1502_5
Chen, H., Liu, Y., Hu, C., Zhang, X., 2023. Vulnerable Road User Trajectory Prediction for Autonomous Driving Using a Data-Driven Integrated Approach. IEEE Trans. Intell. Transp. Syst. 24, 7306–7317. https://doi.org/10.1109/TITS.2023.3254809
Clamann, M., Aubert, M., Cummings, M.L., 2017. Evaluation of Vehicle-to-Pedestrian Communication Displays for Autonomous Vehicles. Presented at the Transportation Research Board 96th Annual MeetingTransportation Research Board.
Cœugnet, S., Cahour, B., Kraiem, S., 2019. Risk-taking, emotions and socio-cognitive dynamics of pedestrian street-crossing decision-making in the city. Transp. Res. Part F Traffic Psychol. Behav. 65, 141–157. https://doi.org/10.1016/j.trf.2019.07.011
Cohen, J., Dearnaley, E.J., Hansel, C.E.M., 1955. The Risk Taken in Crossing a Road. J. Oper. Res. Soc. 6, 120–128. https://doi.org/10.1057/jors.1955.15
Crowley-Koch, B.J., Van Houten, R., Lim, E., 2011. EFFECTS OF PEDESTRIAN PROMPTS ON MOTORIST YIELDING AT

CROSSWALKS. J. Appl. Behav. Anal. 44, 121–126. https://doi.org/10.1901/jaba.2011.44-121
D. Rothenbücher, J. Li, D. Sirkin, B. Mok, W. Ju, 2016. Ghost driver: A field study investigating the interaction between pedestrians and driverless vehicles, in: 2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN). Presented at the 2016 25th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN), pp. 795–802. https://doi.org/10.1109/ROMAN.2016.7745210
de Clercq, K., Dietrich, A., Núñez Velasco, J.P., de Winter, J., Happee, R., 2019. External Human-Machine Interfaces on Automated Vehicles: Effects on Pedestrian Crossing Decisions. Hum. Factors 61, 1353–1370. https://doi.org/10.1177/0018720819836343
Deb, S., Strawderman, L.J., Carruth, D.W., 2018. Investigating pedestrian suggestions for external features on fully autonomous vehicles: A virtual reality experiment. Transp. Res. Part F Traffic Psychol. Behav. 59, 135–149. https://doi.org/10.1016/j.trf.2018.08.016
DeLucia, P.R., 2015. Perception of collision. Camb. Handb. Appl. Percept. Res. Vol I, Cambridge handbooks in psychology. 568–591. https://doi.org/10.1017/CBO9780511973017.035
Deo, N., Trivedi, M.M., 2017. Learning and predicting on-road pedestrian behavior around vehicles, in: 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC). IEEE Press, Yokohama, Japan, pp. 1–6. https://doi.org/10.1109/ITSC.2017.8317865
Department for Transport, 2024. National Travel Survey 2023 technical report.
Department for Transport, 2023. Reported road casualties in Great Britain: pedestrian factsheet, 2023 [WWW Document]. GOV.UK. URL https://www.gov.uk/government/statistics/reported-road-casualties-great-britain-pedestrian-factsheet-2023/reported-road-casualties-in-great-britain-pedestrian-factsheet-2023
Department for Transport, 2015. The Pathway to Driverless Cars Summary report and action plan.
Department of Transport and Main Roads, 2006. Intersections at Grade Chapter 13, Road Planning and Design Manual. Queensland, Australia.
Dey, D., Holländer, K., Berger, M., Eggen, B., Martens, M., Pfleging, B., Terken, J., 2020. Distance-Dependent eHMIs for the Interaction Between Automated Vehicles and Pedestrians, in: 12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI ’20. Association for Computing Machinery, New York, NY, USA, pp. 192–204. https://doi.org/10.1145/3409120.3410642
Dey, D., Martens, M., Eggen, B., Terken, J., 2019. Pedestrian road-crossing willingness as a function of vehicle automation, external appearance, and driving behaviour. Transp. Res. Part F Traffic Psychol. Behav. 65, 191–205. https://doi.org/10.1016/j.trf.2019.07.027
Dey, D., Matviienko, A., Berger, M., Pfleging, B., Martens, M., Terken, J., 2021. Communicating the intention of an automated vehicle to pedestrians: The contributions of eHMI and vehicle behavior. it - Information Technology 63, 123–141. https://doi.org/10.1515/itit-2020-0025
Dey, D., Terken, J., 2017. Pedestrian Interaction with Vehicles: Roles of Explicit and Implicit Communication, in: Proceedings of the 9th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI ’17. Association for Computing Machinery,

New York, NY, USA, pp. 109–113. https://doi.org/10.1145/3122986.3123009
di Pellegrino, G., Fadiga, L., Fogassi, L., Gallese, V., Rizzolatti, G., 1992. Understanding motor events: a neurophysiological study. Exp. Brain Res. 91, 176–180. https://doi.org/10.1007/BF00230027
Dietrich, A., Maruhn, P., Schwarze, L., Bengler, K., 2020. Implicit Communication of Automated Vehicles in Urban Scenarios: Effects of Pitch and Deceleration on Pedestrian Crossing Behavior, in: Ahram, T., Karwowski, W., Pickl, S., Taiar, R. (Eds.), Human Systems Engineering and Design II, Advances in Intelligent Systems and Computing. Springer International Publishing, Cham, pp. 176–181. https://doi.org/10.1007/978-3-030-27928-8_27
Dindo, H., Zambuto, D., Pezzulo, G., 2011. Motor simulation via coupled internal models using sequential Monte Carlo, in: Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume Three, IJCAI'11. AAAI Press, Barcelona, Catalonia, Spain, pp. 2113–2119.
Dipietro, C.M., Dipietro, C.M., King, E.L., 1970. PEDESTRIAN GAP-ACCEPTANCE. Highw. Res. Rec.
Domeyer, J. D. Lee, H. Toyoda, B. Mehler, B. Reimer, 2022. Interdependence in Vehicle-Pedestrian Encounters and its Implications for Vehicle Automation. IEEE Trans. Intell. Transp. Syst. 23, 4122–4134. https://doi.org/10.1109/TITS.2020.3041562
Domeyer, J., Dinparastdjadid, A., Lee, J.D., Douglas, G., Alsaid, A., Price, M., 2019. Proxemics and Kinesics in Automated Vehicle–Pedestrian Communication: Representing Ethnographic Observations. Transp. Res. Rec. 2673, 70–81. https://doi.org/10.1177/0361198119848413
Domeyer, J.E., Lee, J.D., Toyoda, H., 2020. Vehicle Automation–Other Road User Communication and Coordination: Theory and Mechanisms. IEEE Access 8, 19860–19872. https://doi.org/10.1109/ACCESS.2020.2969233
Droit-Volet, S., Delgado, M., Rattat, A.-C., 2006. The development of the ability to judge time in children. Focus Child Psychol. Res. 81–104.
Dytrt, Z., 2023. Research on walking as a mode of transport: Recommendations of the International Walking Data Standard project versus Czech national transport research project Czechia in Motion. Case Stud. Transp. Policy 12, 101003. https://doi.org/10.1016/j.cstp.2023.101003
E. Rehder, H. Kloeden, 2015. Goal-Directed Pedestrian Prediction, in: 2015 IEEE International Conference on Computer Vision Workshop (ICCVW). Presented at the 2015 IEEE International Conference on Computer Vision Workshop (ICCVW), pp. 139–147. https://doi.org/10.1109/ICCVW.2015.28
El Hamdani, S., Benamar, N., Younis, M., 2020. Pedestrian Support in Intelligent Transportation Systems: Challenges, Solutions and Open issues. Transp. Res. Part C Emerg. Technol. 121, 102856. https://doi.org/10.1016/j.trc.2020.102856
Endsley, M.R., 1995. Toward a Theory of Situation Awareness in Dynamic Systems. Hum. Factors 37, 32–64. https://doi.org/10.1518/001872095779049543
Estes, W.K., 1956. The problem of inference from curves based on group data. Psychol. Bull. 53, 134–140. https://doi.org/10.1037/h0045156
European Road Safety Observatory, 2023. European Road Safety Observatory Facts and Figures – Pedestrians - 2023.

Evans, D., Norman, P., 1998. Understanding pedestrians' road crossing decisions: an application of the theory of planned behaviour. Health Educ. Res. 13, 481–489. https://doi.org/10.1093/her/13.4.481-a
Ezzati Amini, R., Katrakazas, C., Antoniou, C., 2019. Negotiation and Decision-Making for a Pedestrian Roadway Crossing: A Literature Review. Sustainability 11. https://doi.org/10.3390/su11236713
F. Camara, N. Bellotto, S. Cosar, F. Weber, D. Nathanael, M. Althoff, J. Wu, J. Ruenz, A. Dietrich, G. Markkula, A. Schieben, F. Tango, N. Merat, C. Fox, 2021. Pedestrian Models for Autonomous Driving Part II: High-Level Models of Human Behavior. IEEE Trans. Intell. Transp. Syst. 22, 5453–5472. https://doi.org/10.1109/TITS.2020.3006767
F. Schneemann, I. Gohl, 2016. Analyzing driver-pedestrian interaction at crosswalks: A contribution to autonomous driving in urban environments, in: 2016 IEEE Intelligent Vehicles Symposium (IV). Presented at the 2016 IEEE Intelligent Vehicles Symposium (IV), pp. 38–43. https://doi.org/10.1109/IVS.2016.7535361
Färber, B., 2016. Communication and Communication Problems Between Autonomous Vehicles and Human Drivers, in: Maurer, M., Gerdes, J.C., Lenz, B., Winner, H. (Eds.), Autonomous Driving: Technical, Legal and Social Aspects. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 125–144. https://doi.org/10.1007/978-3-662-48847-8_7
Faria, J.J., Krause, S., Krause, J., 2010. Collective behavior in road crossing pedestrians: the role of social information. Behav. Ecol. 21, 1236–1242. https://doi.org/10.1093/beheco/arq141
Farrell, S., Lewandowsky, S., 2018. Computational Modeling of Cognition and Behavior. Cambridge University Press, Cambridge. https://doi.org/10.1017/CBO9781316272503
Federal Highway Administration, 2013. FHWA Course on Bicycle and Pedestrian Transportation - Safety.
Feng, Z., Gao, Y., Zhu, D., Chan, H.-Y., Zhao, M., Xue, R., 2024. Impact of risk perception and trust in autonomous vehicles on pedestrian crossing decision: Navigating the social-technological intersection with the ICLV model. Transp. Policy 152, 71–86. https://doi.org/10.1016/j.tranpol.2024.05.001
Ferguson, S., Luders, B., Grande, R.C., How, J.P., 2015. Real-Time Predictive Modeling and Robust Avoidance of Pedestrians with Uncertain, Changing Intentions, in: Akin, H.L., Amato, N.M., Isler, V., van der Stappen, A.F. (Eds.), Algorithmic Foundations of Robotics XI: Selected Contributions of the Eleventh International Workshop on the Algorithmic Foundations of Robotics. Springer International Publishing, Cham, pp. 161–177. https://doi.org/10.1007/978-3-319-16595-0_10
Fetsch, C.R., Turner, A.H., DeAngelis, G.C., Angelaki, D.E., 2009. Dynamic Reweighting of Visual and Vestibular Cues during Self-Motion Perception. J. Neurosci. 29, 15601–15612. https://doi.org/10.1523/JNEUROSCI.2574-09.2009
Fitzpatrick, K., Turner, S., Brewer, M., 2007. Improving pedestrian safety at unsignalized intersections. Ite J. 77, 34–41.
Forrest, A.D., Konca, M., 2007. AUTONOMOUS CARS & SOCIETY. Worcester Polytechnic Institute, Worcester.
Fox, N.A., Bakermans-Kranenburg, M.J., Yoo, K.H., Bowman, L.C., Cannon, E.N., Vanderwert, R.E., Ferrari, P.F., van IJzendoorn, M.H., 2016.

Assessing human mirror activity with EEG mu rhythm: A meta-analysis. Psychol. Bull. 142, 291–313. https://doi.org/10.1037/bul0000031
Friston, K., Frith, C., 2015. A Duet for one. Conscious. Cogn. 36, 390–405. https://doi.org/10.1016/j.concog.2014.12.003
Friston, K., Kiebel, S., 2009a. Predictive coding under the free-energy principle. Philos. Trans. R. Soc. B Biol. Sci. 364, 1211–1221. https://doi.org/10.1098/rstb.2008.0300
Friston, K., Kiebel, S., 2009b. Cortical circuits for perceptual inference. Neural Netw., Cortical Microcircuits 22, 1093–1104. https://doi.org/10.1016/j.neunet.2009.07.023
Friston, K.J., Frith, C.D., 2015. Active inference, communication and hermeneutics. Cortex, Special issue: Prediction in speech and language processing 68, 129–143. https://doi.org/10.1016/j.cortex.2015.03.025
Fuest, T., Michalowski, L., Träris, L., Otto, H., Bengler, K., 2018. Using the Driving Behavior of an Automated Vehicle to Communicate Intentions - A Wizard of Oz Study. https://doi.org/10.1109/ITSC.2018.8569486
Fuller, R., 1984. A conceptualization of driving behaviour as threat avoidance. Ergonomics 27, 1139–1155. https://doi.org/10.1080/00140138408963596
Gandhi, T., Trivedi, M.M., 2008. Computer Vision and Machine Learning for Enhancing Pedestrian Safety, in: Prokhorov, D. (Ed.), Computational Intelligence in Automotive Applications. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 59–77. https://doi.org/10.1007/978-3-540-79257-4_4
Gauvrit, N., Morsanyi, K., 2014. The Equiprobability Bias from a Mathematical and Psychological Perspective. Adv. Cogn. Psychol. 10, 119–130. https://doi.org/10.5709/acp-0163-9
Gawthrop, P., Loram, I., Lakie, M., Gollee, H., 2011. Intermittent control: a computational theory of human control. Biol. Cybern. 104, 31–51. https://doi.org/10.1007/s00422-010-0416-4
Gergely, G., Csibra, G., 2003. Teleological reasoning in infancy: the naïve theory of rational action. Trends Cogn. Sci. 7, 287–292. https://doi.org/10.1016/S1364-6613(03)00128-1
Gergely, G., Nádasdy, Z., Csibra, G., Bíró, S., 1995. Taking the intentional stance at 12 months of age. Cognition 56, 165–193. https://doi.org/10.1016/0010-0277(95)00661-H
Geruschat, D.R., Hassan, S.E., 2005. Driver Behavior in Yielding to Sighted and Blind Pedestrians at Roundabouts. J. Vis. Impair. Blind. 99, 286–302. https://doi.org/10.1177/0145482X0509900504
Gilbert, S.J., Burgess, P.W., 2008. Executive function. Curr. Biol. 18, R110–R114. https://doi.org/10.1016/j.cub.2007.12.014
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 [WWW Document]. Proc. 41st Annu. Meet. Cogn. Sci. Soc. URL https://cognitivesciencesociety.org/cogsci-2019/
Goldhammer, M., Koehler, S., Doll, K., Sick, B., 2015. Camera based pedestrian path prediction by means of polynomial least-squares approximation and multilayer perceptron neural networks. https://doi.org/10.1109/IntelliSys.2015.7361171
Guéguen, N., Meineri, S., Eyssartier, C., 2015. A pedestrian's stare and drivers' stopping behavior: A field experiment at the pedestrian crossing. Saf. Sci. 75, 87–89. https://doi.org/10.1016/j.ssci.2015.01.018

Guest, O., Martin, A.E., 2021. How Computational Modeling Can Force Theory Building in Psychological Science. Perspect. Psychol. Sci. 16, 789–802. https://doi.org/10.1177/1745691620970585
Gutiérrez-Moreno, R., Barea, R., López-Guillén, E., Araluce, J., Bergasa, L.M., 2022. Reinforcement Learning-Based Autonomous Driving at Intersections in CARLA Simulator. Sensors 22, 8373. https://doi.org/10.3390/s22218373
Habibovic, A., Davidsson, J., 2012. Causation mechanisms in car-to-vulnerable road user crashes: Implications for active safety systems. PTW Cogn. Impair. Driv. Saf. 49, 493–500. https://doi.org/10.1016/j.aap.2012.03.022
Habibovic, A., Lundgren, V.M., Andersson, J., Klingegård, M., Lagström, T., Sirkka, A., Fagerlönn, J., Edgren, C., Fredriksson, R., Krupenia, S., Saluäär, D., Larsson, P., 2018. Communicating Intent of Automated Vehicles to Pedestrians. Front. Psychol. 9. https://doi.org/10.3389/fpsyg.2018.01336
Habibovic, A., Tivesten, E., Uchida, N., Bärgman, J., Ljung Aust, M., 2013. Driver behavior in car-to-pedestrian incidents: An application of the Driving Reliability and Error Analysis Method (DREAM). Accid. Anal. Prev. 50, 554–565. https://doi.org/10.1016/j.aap.2012.05.034
Hamed, M.M., 2001. Analysis of pedestrians' behavior at pedestrian crossings. Saf. Sci. 38, 63–82. https://doi.org/10.1016/S0925-7535(00)00058-8
Hariyono, J., Jo, K.-H., 2015. Pedestrian action recognition using motion type classification. https://doi.org/10.1109/CYBConf.2015.7175919
Harrell, W.A., 1993. The Impact of Pedestrian Visibility and Assertiveness on Motorist Yielding. J. Soc. Psychol. 133, 353–360. https://doi.org/10.1080/00224545.1993.9712153
Harrell, W.A., 1991. Factors Influencing Pedestrian Cautiousness in Crossing Streets. J. Soc. Psychol. 131, 367–372. https://doi.org/10.1080/00224545.1991.9713863
Harrell, W.A., Bereska, T., 1992. Gap Acceptance by Pedestrians. Percept. Mot. Skills 75, 432–434. https://doi.org/10.2466/pms.1992.75.2.432
Hashimoto, Y., Gu, Y., Hsu, L.-T., Shunsuke, K., 2015. A Probabilistic Model for the Estimation of Pedestrian Crossing Behavior at Signalized Intersections. https://doi.org/10.1109/ITSC.2015.248
HCM2010, 2010. Highway Capacity Manual 2010 [WWW Document]. URL https://www.trb.org/Publications/Blurbs/164718.aspx
Heathcote, A., Brown, S., Mewhort, D.J.K., 2000. The power law repealed: The case for an exponential law of practice. Psychon. Bull. Rev. 7, 185–207. https://doi.org/10.3758/BF03212979
Helbing, D., Molnár, P., 1995. Social force model for pedestrian dynamics. Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top. 51, 4282–4286. https://doi.org/10.1103/physreve.51.4282
Himanen, V., Kulmala, R., 1988. An application of logit models in analysing the behaviour of pedestrians and car drivers on pedestrian crossings. Accid. Anal. Prev. 20, 187–197. https://doi.org/10.1016/0001-4575(88)90003-6
Hitchcock, A., Mitchell, C.G.B., 1984. Man and his transport behaviour. Part 2a Walking as a means of transport. Transp. Rev. 4, 177–187. https://doi.org/10.1080/01441648408716556
Hoffmann, E.R., 1994. Estimation of time to vehicle arrival: Effects of age on use of available visual information. Perception 23, 947–955. https://doi.org/10.1068/p230947

Holland, C., Hill, R., 2007. The effect of age, gender and driver status on pedestrians' intentions to cross the road in risky situations. Accid. Anal. Prev. 39, 224–237. https://doi.org/10.1016/j.aap.2006.07.003
Holländer, K., Wintersberger, P., Butz, A., 2019. Overtrust in External Cues of Automated Vehicles: An Experimental Investigation, in: Proceedings of the 11th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI '19. Association for Computing Machinery, New York, NY, USA, pp. 211–221. https://doi.org/10.1145/3342197.3344528
Hoogendoorn, S., H.L. Bovy, P., 2003. Simulation of pedestrian flows by optimal control and differential games. Optim. Control Appl. Methods 24, 153–172. https://doi.org/10.1002/oca.727
I. Kotseruba, A. Rasouli, 2023. Intend-Wait-Perceive-Cross: Exploring the Effects of Perceptual Limitations on Pedestrian Decision-Making, in: 2023 IEEE Intelligent Vehicles Symposium (IV). Presented at the 2023 IEEE Intelligent Vehicles Symposium (IV), pp. 1–6. https://doi.org/10.1109/IV55152.2023.10186749
Iacoboni, M., Molnar-Szakacs, I., Gallese, V., Buccino, G., Mazziotta, J.C., Rizzolatti, G., 2005. Grasping the Intentions of Others with One's Own Mirror Neuron System. PLoS Biol. 3, e79. https://doi.org/10.1371/journal.pbio.0030079
Ishaque, M.M., Noland, R.B., 2008. Behavioural Issues in Pedestrian Speed Choice and Street Crossing Behaviour: A Review. Transp. Rev. 28, 61–85. https://doi.org/10.1080/01441640701365239
J. Hariyono, A. Shahbaz, K. -H. Jo, 2015. Estimation of walking direction for pedestrian path prediction from moving vehicle, in: 2015 IEEE/SICE International Symposium on System Integration (SII). Presented at the 2015 IEEE/SICE International Symposium on System Integration (SII), pp. 750–753. https://doi.org/10.1109/SII.2015.7405073
Jara-Ettinger, J., Schulz, L.E., Tenenbaum, J.B., 2020. The Naïve Utility Calculus as a unified, quantitative framework for action understanding. Cognit. Psychol. 123, 101334. https://doi.org/10.1016/j.cogpsych.2020.101334
Jay, M., Régnier, A., Dasnon, A., Brunet, K., Pelé, M., 2020. The light is red: Uncertainty behaviours displayed by pedestrians during illegal road crossing. Accid. Anal. Prev. 135, 105369. https://doi.org/10.1016/j.aap.2019.105369
Jayaraman, S.K., Creech, C., Tilbury, D.M., Yang, X.J., Pradhan, A.K., Tsui, K.M., Robert, L.P., 2019. Pedestrian Trust in Automated Vehicles: Role of Traffic Signal and AV Driving Behavior. Front. Robot. AI 6. https://doi.org/10.3389/frobt.2019.00117
Jeon, H., Lee, S.-H., 2018. From Neurons to Social Beings: Short Review of the Mirror Neuron System Research and Its Socio-Psychological and Psychiatric Implications. Clin. Psychopharmacol. Neurosci. 16, 18–31. https://doi.org/10.9758/cpn.2018.16.1.18
Jiang, X., Wang, W., Mao, Y., Bengler, K., Bubb, H., 2011. Situational Factors of Influencing Drivers to Give Precedence to Jaywalking Pedestrians at Signalized Crosswalk. Int. J. Comput. Intell. Syst. 4, 1407–1414. https://doi.org/10.2991/ijcis.2011.4.6.35
Jin, S., Qu, X., Xu, C., Wang, D.-H., 2013. Dynamic characteristics of traffic flow with consideration of pedestrians' road-crossing behavior. Phys. Stat.

Mech. Its Appl. 392, 3881–3890. https://doi.org/10.1016/j.physa.2013.04.030
Johansson, C., Gårde, P., Leden, L., 2004. The effect of change of code on safety and mobility for children and elderly as pedestrians at marked crosswalks: A case study Comparing Sweden to Finland. Presented at the TRB 2004 Annual Meeting.
K. Tian, G. Markkula, C. Wei, R. Romano, 2020. Creating Kinematics-dependent Pedestrian Crossing Willingness Model When Interacting with Approaching Vehicle, in: 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC). Presented at the 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), pp. 1–6. https://doi.org/10.1109/ITSC45102.2020.9294430
K. Tian, G. Markkula, C. Wei, Y. M. Lee, R. Madigan, T. Hirose, N. Merat, R. Romano, 2024. Deconstructing Pedestrian Crossing Decisions in Interactions With Continuous Traffic: An Anthropomorphic Model. IEEE Trans. Intell. Transp. Syst. 25, 2466–2478. https://doi.org/10.1109/TITS.2023.3323010
Kadali, B.R., Vedagiri, P., 2016. Proactive pedestrian safety evaluation at unprotected mid-block crosswalk locations under mixed traffic conditions. Saf. Sci. 89, 94–105. https://doi.org/10.1016/j.ssci.2016.05.014
Kalatian, A., Farooq, B., 2022. A context-aware pedestrian trajectory prediction framework for automated vehicles. Transp. Res. Part C Emerg. Technol. 134, 103453. https://doi.org/10.1016/j.trc.2021.103453
Kapons, M.M., Kelly, P., 2023. Biased Inference Due to Prior Beliefs: Evidence From the Field. https://doi.org/10.2139/ssrn.4209631
Karasev, V., Ayvaci, A., Heisele, B., Soatto, S., 2016. Intent-aware long-term prediction of pedestrian motion. 2016 IEEE Int. Conf. Robot. Autom. ICRA 2543–2549.
Karras, G.C., Bechlioulis, C.P., Leonetti, M., Palomeras, N., Kormushev, P., Kyriakopoulos, K.J., Caldwell, D.G., 2013. On-line identification of autonomous underwater vehicles through global derivative-free optimization, in: 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems. Presented at the 2013 IEEE/RSJ International Conference on Intelligent Robots and Systems, pp. 3859–3864. https://doi.org/10.1109/IROS.2013.6696908
Katz, A., Zaidel, D., Elgrishi, A., 1975. An Experimental Study of Driver and Pedestrian Interaction during the Crossing Conflict. Hum. Factors 17, 514–527. https://doi.org/10.1177/001872087501700510
Kilner, J.M., 2011. More than one pathway to action understanding. Trends Cogn. Sci. 15, 352–357. https://doi.org/10.1016/j.tics.2011.06.005
Kilner, J.M., Friston, K.J., Frith, C.D., 2007a. Predictive coding: an account of the mirror neuron system. Cogn. Process. 8, 159–166. https://doi.org/10.1007/s10339-007-0170-2
Kilner, J.M., Friston, K.J., Frith, C.D., 2007b. The mirror-neuron system: a Bayesian perspective. Neuroreport 18, 619–623. https://doi.org/10.1097/WNR.0b013e3281139ed0
King, M.J., Soole, D., Ghafourian, A., 2009. Illegal pedestrian crossing at signalised intersections: Incidence and relative risk. Accid. Anal. Prev. 41, 485–490. https://doi.org/10.1016/j.aap.2009.01.008
Kitani, K.M., Ziebart, B.D., Bagnell, J.A., Hebert, M., 2012. Activity Forecasting, in: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (Eds.),

Computer Vision – ECCV 2012. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 201–214.
Koehler, S., 2015. Stereo-Vision-Based Pedestrian's Intention Detection in a Moving Vehicle. https://doi.org/10.1109/ITSC.2015.374
Koehler, S., Goldhammer, M., Bauer, S., Doll, K., Brunsmann, U., Dietmayer, K., 2012. Early detection of the Pedestrian's intention to cross the street, Conference Record - IEEE Conference on Intelligent Transportation Systems. https://doi.org/10.1109/ITSC.2012.6338797
Kohn, S.C., de Visser, E.J., Wiese, E., Lee, Y.-C., Shaw, T.H., 2021. Measurement of Trust in Automation: A Narrative Review and Reference Guide. Front. Psychol. 12. https://doi.org/10.3389/fpsyg.2021.604977
Kollmitz, M., Koller, T., Boedecker, J., Burgard, W., 2020. Learning Human-Aware Robot Navigation from Physical Interaction via Inverse Reinforcement Learning, in: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). Presented at the 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 11025–11031. https://doi.org/10.1109/IROS45743.2020.9340865
Kooij, J.F.P., Schneider, N., Flohr, F., Gavrila, D.M., 2014. Context-Based Pedestrian Path Prediction, in: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (Eds.), Computer Vision – ECCV 2014. Springer International Publishing, Cham, pp. 618–633.
Kooijman, L., Happee, R., de Winter, J.C.F., 2019. How Do eHMIs Affect Pedestrians' Crossing Behavior? A Study Using a Head-Mounted Display Combined with a Motion Suit. Information 10. https://doi.org/10.3390/info10120386
Kudarauskas, N., 2007. Analysis of emergency braking of a vehicle. Transport 22, 154–159. https://doi.org/10.1080/16484142.2007.9638118
Kwak, J.-Y., Ko, B.C., Nam, J.-Y., 2017. Pedestrian intention prediction based on dynamic fuzzy automata for vehicle driving at nighttime. Infrared Phys. Technol. 81, 41–51. https://doi.org/10.1016/j.infrared.2016.12.014
Kwon, M., Biyik, E., Talati, A., Bhasin, K., Losey, D.P., Sadigh, D., 2020. When Humans Aren't Optimal: Robots that Collaborate with Risk-Aware Humans, in: Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction. pp. 43–52. https://doi.org/10.1145/3319502.3374832
L. Crosato, H. P. H. Shum, E. S. L. Ho, C. Wei, 2023. Interaction-Aware Decision-Making for Automated Vehicles Using Social Value Orientation. IEEE Trans. Intell. Veh. 8, 1339–1349. https://doi.org/10.1109/TIV.2022.3189836
L. Griffiths, T., Kemp, C., B. Tenenbaum, J., 2008. Bayesian models of cognition. https://doi.org/10.1184/R1/6613682.v1
Lanzer, M., Babel, F., Yan, F., Zhang, B., You, F., Wang, J., Baumann, M., 2020. Designing Communication Strategies of Autonomous Vehicles with Pedestrians: An Intercultural Study, in: 12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI '20. Association for Computing Machinery, New York, NY, USA, pp. 122–131. https://doi.org/10.1145/3409120.3410653
Lasota, D., Al-Wathinani, A., Krajewski, P., Goniewicz, K., Pawłowski, W., 2020. Alcohol and Road Accidents Involving Pedestrians as Unprotected Road Users. Int. J. Environ. Res. Public. Health 17, 8995. https://doi.org/10.3390/ijerph17238995

Lau, M., Jipp, M., Oehl, M., 2022. Toward a Holistic Communication Approach to an Automated Vehicle's Communication With Pedestrians: Combining Vehicle Kinematics With External Human-Machine Interfaces for Differently Sized Automated Vehicles. Front. Psychol. 13, 882394. https://doi.org/10.3389/fpsyg.2022.882394

Layegh, M., Mirbaha, B., Rassafi, A.A., 2020. Modeling the pedestrian behavior at conflicts with vehicles in multi-lane roundabouts (a cellular automata approach). Phys. Stat. Mech. Its Appl. 556, 124843. https://doi.org/10.1016/j.physa.2020.124843

Lecoutre, M.-P., 1992. Cognitive models and problem spaces in "purely random" situations. Educ. Stud. Math. 23, 557–568. https://doi.org/10.1007/BF00540060

Leden, L., Gårder, P., Johansson, C., 2006. Safe pedestrian crossings for children and elderly. Accid. Anal. Prev. 38, 289–294. https://doi.org/10.1016/j.aap.2005.09.012

Lee, D., Coricelli, G., 2020. An Empirical Test of the Role of Value Certainty in Decision Making. Front. Psychol. 11. https://doi.org/10.3389/fpsyg.2020.574473

Lee, D.G., Hare, T.A., 2023. Value certainty and choice confidence are multidimensional constructs that guide decision-making. Cogn. Affect. Behav. Neurosci. 23, 503–521. https://doi.org/10.3758/s13415-022-01054-4

Lee, D.N., 1976. A Theory of Visual Control of Braking Based on Information about Time-to-Collision. Perception 5, 437–459. https://doi.org/10.1068/p050437

Lee, J.D., Wickens, C.D., Liu, Y., Boyle, L.N., 2017. Designing for People: An introduction to human factors engineering, 3rd ed. CreateSpace Independent Publishing Platform, Charleston, SC.

Lee, M.D., Newell, B.R., 2011. Using hierarchical Bayesian methods to examine the tools of decision-making. Judgm. Decis. Mak. 6, 832–842. https://doi.org/10.1017/S1930297500004253

Lee, S.M., Lim, S., Pathak, R.D., 2011. Culture and entrepreneurial orientation: a multi-country study. Int. Entrep. Manag. J. 7, 1–15. https://doi.org/10.1007/s11365-009-0117-4

Lee, Y.M., Madigan, R., Garcia, J., Tomlinson, A., Solernou, A., Romano, R., Markkula, G., Merat, N., Uttley, J., 2019a. Understanding the Messages Conveyed by Automated Vehicles, in: Proceedings of the 11th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI '19. Association for Computing Machinery, New York, NY, USA, pp. 134–143. https://doi.org/10.1145/3342197.3344546

Lee, Y.M., Madigan, R., Giles, O., Garach-Morcillo, L., Markkula, G., Fox, C., Camara, F., Rothmueller, M., Vendelbo-Larsen, S.A., Rasmussen, P.H., Dietrich, A., Nathanael, D., Portouli, V., Schieben, A., Merat, N., 2021. Road users rarely use explicit communication when interacting in today's traffic: implications for automated vehicles. Cogn. Technol. Work 23, 367–380. https://doi.org/10.1007/s10111-020-00635-y

Lee, Y.M., Madigan, R., Uzondu, C., Garcia, J., Romano, R., Markkula, G., Merat, N., 2022. Learning to interpret novel eHMI: The effect of vehicle kinematics and eHMI familiarity on pedestrian' crossing behavior. J. Safety Res. 80, 270–280. https://doi.org/10.1016/j.jsr.2021.12.010

Lee, Y.M., Sidorov, V., Madigan, R., Garcia de Pedro, J., Markkula, G., Merat, N., 2024. Hello, is it me you're Stopping for? The Effect of external Human Machine Interface Familiarity on Pedestrians' Crossing Behaviour in an Ambiguous Situation. Hum. Factors 00187208241272070. https://doi.org/10.1177/00187208241272070

Lee, Y.M., Uttley, J., Solernou, A., Giles, O.T., Romano, R., Markkula, G., Merat, N., 2019b. Investigating Pedestrians' Crossing Behaviour During Car Deceleration Using Wireless Head Mounted Display: An Application Towards the Evaluation of eHMI of Automated Vehicles. Proc. 10th Int. Driv. Symp. Hum. Factors Driv. Assess. Train. Veh. Des. Driv. Assess. 2019.

Leonetti, M., Ahmadzadeh, S.R., Kormushev, P., 2013. On-line learning to recover from thruster failures on Autonomous Underwater Vehicles, in: 2013 OCEANS - San Diego. Presented at the 2013 OCEANS - San Diego, pp. 1–6. https://doi.org/10.23919/OCEANS.2013.6741265

Leonetti, M., Kormushev, P., Sagratella, S., 2012. Combining Local and Global Direct Derivative-Free Optimization for Reinforcement Learning. Cybern. Inf. Technol. 12, 53–65. https://doi.org/10.2478/cait-2012-0021

Li, T., Hu ,Chengxi, Sze ,N.N., Sun ,Zhanbo, Ding ,Hongliang, and Chen, T., 2025. Modelling pedestrian-vehicle interaction behaviours at non-signalised crosswalks using a game theory approach incorporating the risk perception. Transp. Transp. Sci. 0, 1–31. https://doi.org/10.1080/23249935.2025.2479084

Lin, P.-S., Kourtellis, A., Wang, Z., Chen, C., Rangaswamy, R., Jackman, J., University of South Florida. Center for Urban Transportation Research, 2019. Understanding Interactions between Drivers and Pedestrian Features at Signalized Intersections – Phase 3.

Lin, Y.-S., Srinivasan, A.R., Leonetti, M., Billington, J., Markkula, G., 2022. A Utility Maximization Model of Pedestrian and Driver Interactions. IEEE Access 10, 118888–118899. https://doi.org/10.1109/ACCESS.2022.3213363

Lio, M.D., Dona, R., Papini, G.P.R., Gurney, K., 2020. Agent Architecture for Adaptive Behaviors in Autonomous Driving. IEEE Access 8, 154906–154923. https://doi.org/10.1109/ACCESS.2020.3007018

Liu, H., Hirayama, T., 2025. Pre-instruction for Pedestrians Interacting Autonomous Vehicles with an eHMI: Effects on Their Psychology and Walking Behavior. IEEE Trans. Intell. Transp. Syst. 1–12. https://doi.org/10.1109/TITS.2025.3560621

Ljung Aust, M., Fagerlind, H., Sagberg, F., 2012. Fatal intersection crashes in Norway: Patterns in contributing factors and data collection challenges. Accid. Anal. Prev. 45, 782–791. https://doi.org/10.1016/j.aap.2011.11.001

Lobjois, R., Benguigui, N., Cavallo, V., 2013. The effects of age and traffic density on street-crossing behavior. Accid. Anal. Prev. 53, 166–175. https://doi.org/10.1016/j.aap.2012.12.028

Lobjois, R., Cavallo, V., 2009. The effects of aging on street-crossing behavior: From estimation to actual crossing. Accid. Anal. Prev. 41, 259–267. https://doi.org/10.1016/j.aap.2008.12.001

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, 934–943. https://doi.org/10.1016/j.aap.2006.12.013

Lobo, L., Heras-Escribano, M., Travieso, D., 2018. The History and Philosophy of Ecological Psychology. Front. Psychol. 9. https://doi.org/10.3389/fpsyg.2018.02228
Loukaitou-Sideris, A., 2020. Special issue on walking. Transp. Rev. 40, 131–134. https://doi.org/10.1080/01441647.2020.1712044
Lu, L., Ren, G., Wang, W., Chan, C.-Y., Wang, J., 2016. A cellular automaton simulation model for pedestrian and vehicle interaction behaviors at unsignalized mid-block crosswalks. Accid. Anal. Prev. 95, 425–437. https://doi.org/10.1016/j.aap.2016.04.014
Lucas, C.G., Griffiths, T.L., Xu, F., Fawcett, C., Gopnik, A., Kushnir, T., Markson, L., Hu, J., 2014. The Child as Econometrician: A Rational Model of Preference Understanding in Children. PLOS ONE 9, e92160. https://doi.org/10.1371/journal.pone.0092160
Lyu, W., Cao, Y., Ding, Y., Li, J., Tian, K., Zhang, H., 2024. Not all explicit cues help communicate: Pedestrians' perceptions, fixations, and decisions toward automated vehicles with varied appearance. https://doi.org/10.48550/arXiv.2407.06505
M. Bertozzi, A. Broggi, A. Fascioli, A. Tibaldi, R. Chapuis, F. Chausse, 2004. Pedestrian localization and tracking system with Kalman filtering, in: IEEE Intelligent Vehicles Symposium, 2004. Presented at the IEEE Intelligent Vehicles Symposium, 2004, pp. 584–589. https://doi.org/10.1109/IVS.2004.1336449
Ma, S., Yan, X., Billington, J., Leonetti, M., Merat, N., Markkula, G., 2025. Improving models of pedestrian crossing behavior using neural signatures of decision-making. Transp. Res. Part F Traffic Psychol. Behav. 109, 1491–1506. https://doi.org/10.1016/j.trf.2025.01.047
Mahadevan, K., Somanath, S., Sharlin, E., 2018. Communicating Awareness and Intent in Autonomous Vehicle-Pedestrian Interaction, in: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI '18. Association for Computing Machinery, New York, NY, USA, pp. 1–12. https://doi.org/10.1145/3173574.3174003
Malenje, J.O., Zhao, J., Li, P., Han, Y., 2018. An extended car-following model with the consideration of the illegal pedestrian crossing. Phys. Stat. Mech. Its Appl. 508, 650–661. https://doi.org/10.1016/j.physa.2018.05.074
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 Robot. Autom. Mag. 31, 115–120.
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, pgad163. https://doi.org/10.1093/pnasnexus/pgad163
Markkula, G., Madigan, R., Nathanael, D., Portouli, E., Lee, Y.M., Dietrich, A., Billington, J., Schieben, A., Merat, N., 2020. Defining interactions: a conceptual framework for understanding interactive behaviour in human and automated road traffic. Theor. Issues Ergon. Sci. 21, 728–752. https://doi.org/10.1080/1463922X.2020.1736686
Markkula, G., Romano, R., Madigan, R., Fox, C.W., Giles, O.T., Merat, N., 2018. Models of Human Decision-Making as Tools for Estimating and Optimizing Impacts of Vehicle Automation. Transp. Res. Rec. 2672, 153–163. https://doi.org/10.1177/0361198118792131

Marr, D., 1982. Vision: A computational investigation into the human representation and processing of visual information. W.H. Freeman, San Francisco.
Martinez-Gil, F., Lozano, M., Fernández, F., 2014. MARL-Ped: A multi-agent reinforcement learning based framework to simulate pedestrian groups. Simul. Model. Pract. Theory 47, 259–275. https://doi.org/10.1016/j.simpat.2014.06.005
MATLAB, 2022. anovan [WWW Document]. Anovan N-Way Anal. Var. URL https://www.mathworks.com/help/stats/anovan.html
Meir, A., Parmet, Y., Oron-Gilad, T., 2013. Towards understanding child-pedestrians' hazard perception abilities in a mixed reality dynamic environment. Transp. Res. Part F Traffic Psychol. Behav. 20, 90–107. https://doi.org/10.1016/j.trf.2013.05.004
Merat, N., Louw, T., Madigan, R., Wilbrink, M., Schieben, A., 2018. What externally presented information do VRUs require when interacting with fully Automated Road Transport Systems in shared space? Accid. Anal. Prev. 118, 244–252. https://doi.org/10.1016/j.aap.2018.03.018
Mihet, R., 2013. Effects of culture on firm risk-taking: a cross-country and cross-industry analysis. J. Cult. Econ. 37, 109–151. https://doi.org/10.1007/s10824-012-9186-2
Millard-Ball, A., 2018. Pedestrians, Autonomous Vehicles, and Cities. J. Plan. Educ. Res. 38, 6–12. https://doi.org/10.1177/0739456X16675674
Molenberghs, P., Cunnington, R., Mattingley, J.B., 2012. Brain regions with mirror properties: A meta-analysis of 125 human fMRI studies. Neurosci. Biobehav. Rev. 36, 341–349. https://doi.org/10.1016/j.neubiorev.2011.07.004
Moller, K., Nyberg, T., Tumova, J., Betz, J., 2025. Pedestrian-Aware Motion Planning for Autonomous Driving in Complex Urban Scenarios.
Moore, D., Currano, R., Sirkin, D., Habibovic, A., Lundgren, V.M., Dey, D. (Dave), Holländer, K., 2019. Wizards of WoZ: using controlled and field studies to evaluate AV-pedestrian interactions, in: Proceedings of the 11th International Conference on Automotive User Interfaces and Interactive Vehicular Applications: Adjunct Proceedings, AutomotiveUI '19. Association for Computing Machinery, New York, NY, USA, pp. 45–49. https://doi.org/10.1145/3349263.3350756
Moore, R.L., 1953. Pedestrian Choice and Judgment. J. Oper. Res. Soc. 4, 3–10. https://doi.org/10.1057/jors.1953.2
Moussaïd, M., Helbing, D., Theraulaz, G., 2011. How simple rules determine pedestrian behavior and crowd disasters. Proc. Natl. Acad. Sci. 108, 6884–6888. https://doi.org/10.1073/pnas.1016507108
Moussaïd, M., Perozo, N., Garnier, S., Helbing, D., Theraulaz, G., 2010. The walking behaviour of pedestrian social groups and its impact on crowd dynamics. PloS One 5, e10047. https://doi.org/10.1371/journal.pone.0010047
Myers, C.E., Interian, A., Moustafa, A.A., 2022. A practical introduction to using the drift diffusion model of decision-making in cognitive psychology, neuroscience, and health sciences. Front. Psychol. 13. https://doi.org/10.3389/fpsyg.2022.1039172
Nasernejad, P., Sayed, T., Alsaleh, R., 2023. Multiagent modeling of pedestrian-vehicle conflicts using Adversarial Inverse Reinforcement Learning. Transp. Transp. Sci. 19, 2061081. https://doi.org/10.1080/23249935.2022.2061081

National Highway Traffic Safety Administration, 2024. NHTSA Launches Put the Phone Away or Pay Campaign; Releases 2023 Fatality Early Estimates | NHTSA [WWW Document]. URL https://www.nhtsa.gov/press-releases/2022-traffic-deaths-2023-early-estimates
Naumann, R.B., 2025. Pedestrian and Overall Road Traffic Crash Deaths — United States and 27 Other High-Income Countries, 2013–2022. MMWR Morb. Mortal. Wkly. Rep. 74. https://doi.org/10.15585/mmwr.mm7408a2
Nissan Motor Corporation, 2015. Nissan IDS Concept: Nissan's vision for the future of EVs and autonomous driving [WWW Document]. Off. Eur. Newsroom. URL https://europe.nissannews.com/en-GB/releases/nissan-ids-concept-nissan-s-vision-for-the-future-of-evs-and-autonomous-driving
Nissan Motor Corporation, 2013. Nissan Announces Unprecedented Autonomous Drive Benchmarks.
Niv, Y., 2021. The primacy of behavioral research for understanding the brain. Behav. Neurosci. 135, 601–609. https://doi.org/10.1037/bne0000471
Nordhoff, S., Hagenzieker, M., Lee, Y.M., Wilbrink, M., Merat, N., Oehl, M., 2025. "It's just another car driving" − Perceptions of U.S. residents interacting with driverless automated vehicles on public roads. Transp. Res. Part F Traffic Psychol. Behav. 111, 188–210. https://doi.org/10.1016/j.trf.2025.01.024
Nuñez Velasco, J.P., Farah, H., van Arem, B., Hagenzieker, M.P., 2019. Studying pedestrians' crossing behavior when interacting with automated vehicles using virtual reality. Transp. Res. Part F Traffic Psychol. Behav. 66, 1–14. https://doi.org/10.1016/j.trf.2019.08.015
Oberman, L.M., Pineda, J.A., Ramachandran, V.S., 2007. The human mirror neuron system: a link between action observation and social skills. Soc. Cogn. Affect. Neurosci. 2, 62–66. https://doi.org/10.1093/scan/nsl022
O'Dowd, E., Pollet, T.V., 2018. Gender Differences in Use of a Pedestrian Crossing: An Observational Study in Newcastle upon Tyne. Lett. Evol. Behav. Sci. 9, 1–4. https://doi.org/10.5178/lebs.2018.65
Orban, G.A., Lanzilotto, M., Bonini, L., 2021. From Observed Action Identity to Social Affordances. Trends Cogn. Sci. 25, 493–505. https://doi.org/10.1016/j.tics.2021.02.012
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, 962–971. https://doi.org/10.1016/j.aap.2005.04.017
Pallier, G., Wilkinson, R., Danthiir, V., Kleitman, S., Knezevic, G., Stankov, L., Roberts, R.D., 2002. The role of individual differences in the accuracy of confidence judgments. J. Gen. Psychol. 129, 257–299. https://doi.org/10.1080/00221300209602099
Palmeiro, A., van der Kint, S., Vissers, L., Farah, H., de Winter, J.C.F., Hagenzieker, M., 2018. Interaction between pedestrians and automated vehicles: A Wizard of Oz experiment. Transp. Res. Part F Traffic Psychol. Behav. 58, 1005–1020. https://doi.org/10.1016/j.trf.2018.07.020
Papadimitriou, E., Lassarre, S., Yannis, G., 2016. Pedestrian Risk Taking While Road Crossing: A Comparison of Observed and Declared Behaviour. Transp. Res. Arena TRA2016 14, 4354–4363. https://doi.org/10.1016/j.trpro.2016.05.357
Papadimitriou, E., Yannis, G., Golias, J., 2009. A critical assessment of pedestrian behaviour models. Transp. Res. Part F Traffic Psychol. Behav. 12, 242–255. https://doi.org/10.1016/j.trf.2008.12.004

Park, H., Oh, T., Kim, I., 2024. Effects of driver's braking behavior by the real-time pedestrian scale warning system. Accid. Anal. Prev. 205, 107685. https://doi.org/10.1016/j.aap.2024.107685
Pawar, D., Kumar, V., Singh, N., Patil, G., 2016. Analysis of dilemma zone for pedestrians at high-speed uncontrolled midblock crossing. Transp. Res. Part C Emerg. Technol. 70. https://doi.org/10.1016/j.trc.2016.04.012
Pawar, D.S., Patil, G.R., 2016. Critical gap estimation for pedestrians at uncontrolled mid-block crossings on high-speed arterials. Saf. Sci. 86, 295–303. https://doi.org/10.1016/j.ssci.2016.03.011
Pawar, D.S., Patil, G.R., 2015. Pedestrian temporal and spatial gap acceptance at mid-block street crossing in developing world. J. Safety Res. 52, 39–46. https://doi.org/10.1016/j.jsr.2014.12.006
Pekkanen, J., Giles, O.T., Lee, Y.M., Madigan, R., Daimon, T., Merat, N., Markkula, G., 2022. Variable-Drift Diffusion Models of Pedestrian Road-Crossing Decisions. Comput. Brain Behav. 5, 60–80. https://doi.org/10.1007/s42113-021-00116-z
Petzoldt, T., 2016. Size speed bias or size arrival effect—How judgments of vehicles' approach speed and time to arrival are influenced by the vehicles' size. Accid. Anal. Prev. 95, 132–137. https://doi.org/10.1016/j.aap.2016.07.010
Petzoldt, T., 2014. On the relationship between pedestrian gap acceptance and time to arrival estimates. Accid. Anal. Prev. 72, 127–133. https://doi.org/10.1016/j.aap.2014.06.019
Petzoldt, T., Schleinitz, K., Banse, R., 2018. Potential safety effects of a frontal brake light for motor vehicles. IET Intell. Transp. Syst. 12, 449–453. https://doi.org/10.1049/iet-its.2017.0321
Pezzulo, G., Donnarumma, F., Dindo, H., 2013. Human Sensorimotor Communication: A Theory of Signaling in Online Social Interactions. PLOS ONE 8, e79876. https://doi.org/10.1371/journal.pone.0079876
Phillips, D.P., Brewer, K.M., 2011. The relationship between serious injury and blood alcohol concentration (BAC) in fatal motor vehicle accidents: BAC = 0.01% is associated with significantly more dangerous accidents than BAC = 0.00%. Addiction 106, 1614–1622. https://doi.org/10.1111/j.1360-0443.2011.03472.x
Piaget, J., 1970. Child's Conception of Movement and Speed, Routledge Library Editions: Piaget. ed, Piaget. Routledge.
Pikūnas, A., Pumputis ,Vidmantas, and Sadauskas, V., 2004. The influence of vehicles speed on accident rates and their consequences. Transport 19, 15–19. https://doi.org/10.1080/16484142.2004.9637946
Prédhumeau, M., Mancheva, L., Dugdale, J., Spalanzani, A., 2022. Agent-Based Modeling for Predicting Pedestrian Trajectories Around an Autonomous Vehicle. J. Artif. Intell. Res. 73, 1385–1433. https://doi.org/10.1613/jair.1.13425
Price, W.L., 1983. Global optimization by controlled random search. J. Optim. Theory Appl. 40, 333–348. https://doi.org/10.1007/BF00933504
Proverbio, A.M., Zani, A., 2023. Mirror Neurons in Action: ERPs and Neuroimaging Evidence, in: Boggio, P.S., Wingenbach, T.S.H., da Silveira Coêlho, M.L., Comfort, W.E., Murrins Marques, L., Alves, M.V.C. (Eds.), Social and Affective Neuroscience of Everyday Human Interaction: From Theory to Methodology. Springer, Cham (CH).
Pugliese, B.J., Barton, B.K., Davis, S.J., Lopez, G., 2020. Assessing pedestrian safety across modalities via a simulated vehicle time-to-arrival task.

Accid. Anal. Prev. 134, 105344. https://doi.org/10.1016/j.aap.2019.105344
Quintero, R., Parra, I., Llorca, D.F., Sotelo, M.Á., 2015. Pedestrian Intention and Pose Prediction through Dynamical Models and Behaviour Classification. 2015 IEEE 18th Int. Conf. Intell. Transp. Syst. 83–88.
R. Quintero, I. Parra, D. F. Llorca, M. A. Sotelo, 2014. Pedestrian path prediction based on body language and action classification, in: 17th International IEEE Conference on Intelligent Transportation Systems (ITSC). Presented at the 17th International IEEE Conference on Intelligent Transportation Systems (ITSC), pp. 679–684. https://doi.org/10.1109/ITSC.2014.6957768
Raff, M.S., Hart, J.W., 1950. A Volume Warrant for Urban Stop Signs.
Raghuram Kadali, B., Rathi, N., Perumal, V., 2014. Evaluation of pedestrian mid-block road crossing behaviour using artificial neural network. J. Traffic Transp. Eng. Engl. Ed. 1, 111–119. https://doi.org/10.1016/S2095-7564(15)30095-7
Ramírez, M., Geffner, H., 2010. Probabilistic Plan Recognition Using Off-the-Shelf Classical Planners. Proc. AAAI Conf. Artif. Intell. 24, 1121–1126. https://doi.org/10.1609/aaai.v24i1.7745
Rampf, F., Grigoropoulos, G., Malcolm, P., Keler, A., Bogenberger, K., 2023. Modelling autonomous vehicle interactions with bicycles in traffic simulation. Front. Future Transp. 3. https://doi.org/10.3389/ffutr.2022.894148
Ranga, A., Giruzzi, F., Bhanushali, J., Wirbel, E., Pérez, P., Vu, T.-H., Perrotton, X., 2020. VRUNet: Multi-Task Learning Model for Intent Prediction of Vulnerable Road Users. Electron. Imaging 32, 109-1-109–10. https://doi.org/10.2352/ISSN.2470-1173.2020.16.AVM-109
Räsänen, M., Summala, H., 1998. Attention and expectation problems in bicycle–car collisions: an in-depth study. Accid. Anal. Prev. 30, 657–666. https://doi.org/10.1016/S0001-4575(98)00007-4
Rasouli, A., Kotseruba, I., Tsotsos, J.K., 2018. Understanding Pedestrian Behavior in Complex Traffic Scenes. IEEE Trans. Intell. Veh. 3, 61–70. https://doi.org/10.1109/TIV.2017.2788193
Ratcliff, R., Smith, P.L., Brown, S.D., McKoon, G., 2016. Diffusion Decision Model: Current Issues and History. Trends Cogn. Sci. 20, 260–281. https://doi.org/10.1016/j.tics.2016.01.007
Razmi Rad, S., Homem de Almeida Correia, G., Hagenzieker, M., 2020. Pedestrians' road crossing behaviour in front of automated vehicles: Results from a pedestrian simulation experiment using agent-based modelling. Transp. Res. Part F Traffic Psychol. Behav. 69, 101–119. https://doi.org/10.1016/j.trf.2020.01.014
Reimer, B., 2014. Driver Assistance Systems and the Transition to Automated Vehicles: A Path to Increase Older Adult Safety and Mobility? Public Policy Aging Rep. 24, 27–31. https://doi.org/10.1093/ppar/prt006
Rezwana, S., Lownes, N., 2024. Interactions and Behaviors of Pedestrians with Autonomous Vehicles: A Synthesis. Future Transp. 4, 722–745. https://doi.org/10.3390/futuretransp4030034
Rezwana, S., Shaon, M.R.R., Lownes, N., Jackson, E., 2025. Bridging the gap: understanding the factors affecting pedestrian safety perceptions in the age of driverless vehicles. Traffic Saf. Res. 9, e000080–e000080. https://doi.org/10.55329/pjax7195

Ribino, P., 2023. The role of politeness in human–machine interactions: a systematic literature review and future perspectives. Artif. Intell. Rev. 56, 445–482. https://doi.org/10.1007/s10462-023-10540-1
Risser, R., 1985. Behavior in traffic conflict situations. Accid. Anal. Prev. 17, 179–197. https://doi.org/10.1016/0001-4575(85)90020-X
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. University of Iowa. https://doi.org/10.17077/drivingassessment.1633
Rizzolatti, G., Craighero, L., 2004. The Mirror-Neuron System. Annu. Rev. Neurosci. 27, 169–192. https://doi.org/10.1146/annurev.neuro.27.070203.144230
Rizzolatti, G., Fogassi, L., Gallese, V., 2001. Neurophysiological mechanisms underlying the understanding and imitation of action. Nat. Rev. Neurosci. 2, 661–670. https://doi.org/10.1038/35090060
Rizzolatti, G., Sinigaglia, C., 2016. The mirror mechanism: a basic principle of brain function. Nat. Rev. Neurosci. 17, 757–765. https://doi.org/10.1038/nrn.2016.135
Ropaka, M., Nikolaou, D., Yannis, G., 2020. Investigation of traffic and safety behavior of pedestrians while texting or web-surfing. Traffic Inj. Prev. 21, 389–394. https://doi.org/10.1080/15389588.2020.1770741
Rosenbloom, T., 2009. Crossing at a red light: Behaviour of individuals and groups. Transp. Res. Part F Traffic Psychol. Behav. 12, 389–394. https://doi.org/10.1016/j.trf.2009.05.002
Rosenbloom, T., Ben-Eliyahu, A., Nemrodov, D., 2008. Children's crossing behavior with an accompanying adult. Saf. Sci. 46, 1248–1254. https://doi.org/10.1016/j.ssci.2007.07.004
Rudin, C., 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat. Mach. Intell. 1, 206–215. https://doi.org/10.1038/s42256-019-0048-x
Sadigh, D., Landolfi, N., Sastry, S.S., Seshia, S.A., Dragan, A.D., 2018. Planning for cars that coordinate with people: leveraging effects on human actions for planning and active information gathering over human internal state. Auton. Robots 42, 1405–1426. https://doi.org/10.1007/s10514-018-9746-1
SAE International, 2021. Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles J3016_202104.
Saleh, K., Hossny, M., Nahavandi, S., 2020. Contextual Recurrent Predictive Model for Long-Term Intent Prediction of Vulnerable Road Users. IEEE Trans. Intell. Transp. Syst. 21, 3398–3408. https://doi.org/10.1109/TITS.2019.2927770
Saleh, K., Hossny, M., Nahavandi, S., 2018a. Intent Prediction of Pedestrians via Motion Trajectories Using Stacked Recurrent Neural Networks. IEEE Trans. Intell. Veh. 3, 414–424. https://doi.org/10.1109/TIV.2018.2873901
Saleh, K., Hossny, M., Nahavandi, S., 2018b. Long-Term Recurrent Predictive Model for Intent Prediction of Pedestrians via Inverse Reinforcement Learning, in: 2018 Digital Image Computing: Techniques and Applications (DICTA). Presented at the 2018 Digital Image Computing: Techniques and Applications (DICTA), pp. 1–8. https://doi.org/10.1109/DICTA.2018.8615854
Saleh, K., Hossny, M., Nahavandi, S., 2017a. Towards trusted autonomous vehicles from vulnerable road users perspective, in: 2017 Annual IEEE

International Systems Conference (SysCon). Presented at the 2017 Annual IEEE International Systems Conference (SysCon), pp. 1–7. https://doi.org/10.1109/SYSCON.2017.7934782
Saleh, K., Hossny, M., Nahavandi, S., 2017b. Early intent prediction of vulnerable road users from visual attributes using multi-task learning network, in: 2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC). Presented at the 2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 3367–3372. https://doi.org/10.1109/SMC.2017.8123150
Saleh, K., Hossny, M., Nahavandi, S., 2017c. Intent prediction of vulnerable road users from motion trajectories using stacked LSTM network, in: 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC). Presented at the 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC), pp. 327–332. https://doi.org/10.1109/ITSC.2017.8317941
Scheller, M., Aston, S., Slater, H., Nardini, M., 2025. Learning new perceptual skills: Individual differences in the computations that integrate novel sensory cues into depth perception. https://doi.org/10.1101/2025.03.25.645045
Schieben, A., Wilbrink, M., Kettwich, C., Madigan, R., Louw, T., Merat, N., 2019. Designing the interaction of automated vehicles with other traffic participants: design considerations based on human needs and expectations. Cogn. Technol. Work 21, 69–85. https://doi.org/10.1007/s10111-018-0521-z
Schiff, W., Oldak, R., 1990. Accuracy of judging time to arrival: Effects of modality, trajectory, and gender. J. Exp. Psychol. Hum. Percept. Perform. 16, 303–316. https://doi.org/10.1037/0096-1523.16.2.303
Schmidt, H., Terwilliger, J., AlAdawy, D., Fridman, L., 2019. Hacking Nonverbal Communication Between Pedestrians and Vehicles in Virtual Reality. https://doi.org/10.48550/arXiv.1904.01931
Schmidt, S., Färber, B., 2009. Pedestrians at the kerb – Recognising the action intentions of humans. Transp. Res. Part F Traffic Psychol. Behav. 12, 300–310. https://doi.org/10.1016/j.trf.2009.02.003
Schneider, N., Gavrila, D.M., 2013. Pedestrian Path Prediction with Recursive Bayesian Filters: A Comparative Study, in: Weickert, J., Hein, M., Schiele, B. (Eds.), Pattern Recognition. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 174–183.
Schooler, L., Okhan, M., Hollander, S., Gill, M., Zoh, Y., Crockett, M.J., Yu, H., 2024. Confidence in Moral Decision-Making. Collabra Psychol. 10, 121387. https://doi.org/10.1525/collabra.121387
Schroeder, B.J., 2008. A behavior-based methodology for evaluating pedestrian-vehicle interaction at crosswalks. North Carolina State University.
Schwarting, W., Pierson, A., Alonso-Mora, J., Karaman, S., Rus, D., 2019. Social behavior for autonomous vehicles. Proc. Natl. Acad. Sci. 116, 24972–24978. https://doi.org/10.1073/pnas.1820676116
Sharma, N., Dhiman, C., Indu, S., 2022. Pedestrian Intention Prediction for Autonomous Vehicles: A Comprehensive Survey. Neurocomputing 508, 120–152. https://doi.org/10.1016/j.neucom.2022.07.085
Shepard, R.N., 1987. Toward a Universal Law of Generalization for Psychological Science. Science 237, 1317–1323. https://doi.org/10.1126/science.3629243

Sheykhfard, A., Haghighi, F., 2018. Behavioral analysis of vehicle-pedestrian interactions in Iran. Sci. Iran. 25, 1968–1976. https://doi.org/10.24200/sci.2017.4201
Soares, F., Silva, E., Pereira, F., Silva, C., Sousa, E., Freitas, E., 2021. To cross or not to cross: Impact of visual and auditory cues on pedestrians' crossing decision-making. Transp. Res. Part F Traffic Psychol. Behav. 82, 202–220. https://doi.org/10.1016/j.trf.2021.08.014
Song, X., Han, D., Sun, J., Zhang, Z., 2018. A data-driven neural network approach to simulate pedestrian movement. Phys. Stat. Mech. Its Appl. 509, 827–844. https://doi.org/10.1016/j.physa.2018.06.045
Sripada, A., Bazilinskyy ,Pavlo, and de Winter, J., 2021. Automated vehicles that communicate implicitly: examining the use of lateral position within the lane. Ergonomics 64, 1416–1428. https://doi.org/10.1080/00140139.2021.1925353
Šrol, J., De Neys, W., 2021. Predicting individual differences in conflict detection and bias susceptibility during reasoning. Think. Reason. 27, 38–68. https://doi.org/10.1080/13546783.2019.1708793
Staplin, L., Lococo, K., Byington, S., Harkey, D., 2001. Guidelines And Recommendations To Accommodate Older Drivers and Pedestrians, May 2001 - FHWA-RD-01-051 (No. FHWA–RD–01–051).
Sucha, M., Dostal, D., Risser, R., 2017. Pedestrian-driver communication and decision strategies at marked crossings. Accid. Anal. Prev. 102, 41–50. https://doi.org/10.1016/j.aap.2017.02.018
Sueur, C., Class, B., Hamm, C., Meyer, X., Pelé, M., 2013. Different risk thresholds in pedestrian road crossing behaviour: A comparison of French and Japanese approaches. Accid. Anal. Prev. 58, 59–63. https://doi.org/10.1016/j.aap.2013.04.027
Sun, D., Ukkusuri, S., Benekohal, R., Waller, S., 2002. Modeling of motorist-pedestrian interaction at uncontrolled mid-block crosswalks. Urbana 51.
T. Wang, J. Wu, P. Zheng, M. McDonald, 2010. Study of pedestrians' gap acceptance behavior when they jaywalk outside crossing facilities, in: 13th International IEEE Conference on Intelligent Transportation Systems. Presented at the 13th International IEEE Conference on Intelligent Transportation Systems, pp. 1295–1300. https://doi.org/10.1109/ITSC.2010.5625157
Tabone, W., Lee, Y.M., Merat, N., Happee, R., de Winter, J., 2021. Towards future pedestrian-vehicle interactions: Introducing theoretically-supported AR prototypes, in: 13th International Conference on Automotive User Interfaces and Interactive Vehicular Applications, AutomotiveUI '21. Association for Computing Machinery, New York, NY, USA, pp. 209–218. https://doi.org/10.1145/3409118.3475149
Tefft, B.C., 2013. Impact speed and a pedestrian's risk of severe injury or death. Accid. Anal. Prev. 50, 871–878. https://doi.org/10.1016/j.aap.2012.07.022
The Highway Code [WWW Document], 2023. . GOV.UK. URL https://www.gov.uk/guidance/the-highway-code (accessed 3.9.25).
Thomas, P., Morris, A., Talbot, R., Fagerlind, H., 2013. Identifying the causes of road crashes in Europe. Ann. Adv. Automot. Med. 57, 13–22.
Tian, K., 2023. Psychological Mechanisms in Pedestrian Road Crossing Behaviour: Observations and Models (phd). University of Leeds.
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. https://doi.org/10.1016/j.ssci.2022.105837
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. https://doi.org/10.1016/j.aap.2023.107173
Tian, K., Wei, C., Lyu, W., Wang, Y., Lee, Y.M., Merat, N., Romano, R., Markkula, G., 2025. Interacting With Yielding Vehicles: A Perceptually Plausible Model for Pedestrian Road Crossing Decisions. IEEE Trans. Intell. Transp. Syst. 1–15. https://doi.org/10.1109/TITS.2025.3562899
Tiwari, G., Bangdiwala, S., Saraswat, A., Gaurav, S., 2007. Survival analysis: Pedestrian risk exposure at signalized intersections. Transp. Res. Part F Traffic Psychol. Behav. 10, 77–89. https://doi.org/10.1016/j.trf.2006.06.002
Tsui, K.M., Desai, M., Yanco, H.A., 2010. Considering the bystander's perspective for indirect human-robot interaction, in: 2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI). Presented at the 2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI), pp. 129–130. https://doi.org/10.1109/HRI.2010.5453230
Tump, A.N., Pleskac, T.J., Kurvers, R.H.J.M., 2020. Wise or mad crowds? The cognitive mechanisms underlying information cascades. Sci. Adv. 6, eabb0266. https://doi.org/10.1126/sciadv.abb0266
Turner, S., Fitzpatrick, K., Brewer, M., Park, E.S., 2006. Motorist Yielding to Pedestrians at Unsignalized Intersections: Findings from a National Study on Improving Pedestrian Safety. Transp. Res. Rec. 1982, 1–12. https://doi.org/10.1177/0361198106198200102
Tversky, A., Kahneman, D., 1974. Judgment under Uncertainty: Heuristics and Biases. Science 185, 1124–1131. https://doi.org/10.1126/science.185.4157.1124
Unsworth, N., Brewer ,Gene A., and Spillers, G.J., 2011. Inter- and intra-individual variation in immediate free recall: An examination of serial position functions and recall initiation strategies. Memory 19, 67–82. https://doi.org/10.1080/09658211.2010.535658
Uono, S., Hietanen, J.K., 2015. Eye Contact Perception in the West and East: A Cross-Cultural Study. PLoS ONE 10.
Valos, N., Bennett, J.M., 2023. The relationship between cognitive functioning and street-crossing behaviours in adults: A systematic review and meta-analysis. Transp. Res. Part F Traffic Psychol. Behav. 99, 356–373. https://doi.org/10.1016/j.trf.2023.10.018
Vandekerckhove, J., Tuerlinckx, F., Lee, M.D., 2008. A Bayesian Approach to Diffusion Process Models of Decision-Making.
Várhelyi, A., 1998. Drivers' speed behaviour at a zebra crossing: a case study. Accid. Anal. Prev. 30, 731–743. https://doi.org/10.1016/S0001-4575(98)00026-8
Vilares, I., Kording, K., 2011. Bayesian models: the structure of the world, uncertainty, behavior, and the brain. Ann. N. Y. Acad. Sci. 1224, 22–39. https://doi.org/10.1111/j.1749-6632.2011.05965.x
Wang, M., Hoogendoorn, S.P., Daamen, W., van Arem, B., Happee, R., 2015. Game theoretic approach for predictive lane-changing and car-following

control. Transp. Res. Part C Emerg. Technol. 58, 73–92. https://doi.org/10.1016/j.trc.2015.07.009
Wang, Y., Srinivasan, A.R., Jokinen, J.P.P., Oulasvirta, A., Markkula, G., 2025. Pedestrian crossing decisions can be explained by bounded optimal decision-making under noisy visual perception. Transp. Res. Part C Emerg. Technol. 171, 104963. https://doi.org/10.1016/j.trc.2024.104963
Wang, Z., Zheng, O., Li, L., Abdel-Aty, M., Cruz-Neira, C., Islam, Z., 2023. Towards Next Generation of Pedestrian and Connected Vehicle In-the-loop Research: A Digital Twin Co-Simulation Framework. IEEE Trans. Intell. Veh. 8, 2674–2683. https://doi.org/10.1109/TIV.2023.3250353
Wann, J.P., Poulter, D.R., Purcell, C., 2011. Reduced sensitivity to visual looming inflates the risk posed by speeding vehicles when children try to cross the road. Psychol. Sci. 22, 429–434. https://doi.org/10.1177/0956797611400917
Wessels, M., Zähme, C., Oberfeld, D., 2023. Auditory Information Improves Time-to-collision Estimation for Accelerating Vehicles. Curr. Psychol. 42, 23195–23205. https://doi.org/10.1007/s12144-022-03375-6
Whetsel Borzendowski, S.A., Rosenberg, R.L., Sewall, A.S., Tyrrell, R.A., 2013. Pedestrians' estimates of their own nighttime conspicuity are unaffected by severe reductions in headlight illumination. J. Safety Res. 47, 25–30. https://doi.org/10.1016/j.jsr.2013.08.007
Whiten, A., 1991. Natural Theories of Mind: Evolution, Development and Simulation of Everyday Mindreading. Blackwell.
Wigan, M., 1995. Treatment of walking as a mode of transportation.
Wilkinson, G.N., Rogers, C.E., 1973. Symbolic Description of Factorial Models for Analysis of Variance. J. R. Stat. Soc. Ser. C Appl. Stat. 22, 392–399. https://doi.org/10.2307/2346786
Willis, A., Gjersoe, N., Havard, C., Kerridge, J., Kukla, R., 2004. Human Movement Behaviour in Urban Spaces: Implications for the Design and Modelling of Effective Pedestrian Environments. Environ. Plan. B Plan. Des. 31, 805–828. https://doi.org/10.1068/b3060
Woodward, A.L., 1998. Infants selectively encode the goal object of an actor's reach. Cognition 69, 1–34. https://doi.org/10.1016/S0010-0277(98)00058-4
World Health Organization, 2023a. Road traffic injuries.
World Health Organization, 2023b. Global status report on road safety 2023. World Health Organization.
World Health Organization, 2013. Pedestrian safety: a road safety manual for decision-makers and practitioners. World Health Organization.
World Health Organization, 2004. World report on road traffic injury prevention.
Wright, J.R., Leyton-Brown, K., 2017. Predicting human behavior in unrepeated, simultaneous-move games. Games Econ. Behav. 106, 16–37. https://doi.org/10.1016/j.geb.2017.09.009
Wu, W., Chen, R., Jia, H., Li, Y., Liang, Z., 2019. Game theory modeling for vehicle–pedestrian interactions and simulation based on cellular automata. Int. J. Mod. Phys. C 30, 1950025. https://doi.org/10.1142/S0129183119500256
Wynne, R.A., Beanland, V., Salmon, P.M., 2019. Systematic review of driving simulator validation studies. Saf. Sci. 117, 138–151. https://doi.org/10.1016/j.ssci.2019.04.004
Y. Hashimoto, Y. Gu, L. -T. Hsu, S. Kamijo, 2015. Probability estimation for pedestrian crossing intention at signalized crosswalks, in: 2015 IEEE

International Conference on Vehicular Electronics and Safety (ICVES). Presented at the 2015 IEEE International Conference on Vehicular Electronics and Safety (ICVES), pp. 114–119. https://doi.org/10.1109/ICVES.2015.7396904
Y. Ma, E. W. M. Lee, R. K. K. Yuen, 2016. An Artificial Intelligence-Based Approach for Simulating Pedestrian Movement. IEEE Trans. Intell. Transp. Syst. 17, 3159–3170. https://doi.org/10.1109/TITS.2016.2542843
Y. Wang, A. Ramakrishnan Srinivasan, J. P. P. Jokinen, A. Oulasvirta, G. Markkula, 2023. Modeling human road crossing decisions as reward maximization with visual perception limitations, in: 2023 IEEE Intelligent Vehicles Symposium (IV). Presented at the 2023 IEEE Intelligent Vehicles Symposium (IV), pp. 1–6. https://doi.org/10.1109/IV55152.2023.10186617
Yang, Y., Lee, Y.M., Kalantari, A.H., de Pedro, J.G., Horrobin, A., Daly, M., Solernou, A., Holmes, C., Markkula, G., Merat, N., 2024. Using distributed simulations to investigate driver-pedestrian interactions and kinematic cues: Implications for automated vehicle behaviour and communication. Transp. Res. Part F Traffic Psychol. Behav. 107, 84–97. https://doi.org/10.1016/j.trf.2024.08.027
Yannis, G., Nikolaou, D., Laiou, A., Stürmer, Y.A., Buttler, I., Jankowska-Karpa, D., 2020. Vulnerable road users: Cross-cultural perspectives on performance and attitudes. IATSS Res. 44, 220–229. https://doi.org/10.1016/j.iatssr.2020.08.006
Yannis, G., Papadimitriou ,E., and Theofilatos, A., 2013. Pedestrian gap acceptance for mid-block street crossing. Transp. Plan. Technol. 36, 450–462. https://doi.org/10.1080/03081060.2013.818274
Yi-Rong, K., Di-Hua, S., Shu-Hong, Y., 2015. A new car-following model considering driver's individual anticipation behavior. Nonlinear Dyn. 82, 1293–1302. https://doi.org/10.1007/s11071-015-2236-5
Yue, L., Abdel-Aty, M., Wu, Y., Zheng, O., Yuan, J., 2020. In-depth approach for identifying crash causation patterns and its implications for pedestrian crash prevention. J. Safety Res. 73, 119–132. https://doi.org/10.1016/j.jsr.2020.02.020
Zegeer, C.V., Bushell, M., 2006. Pedestrian crash trends and potential countermeasures from around the world. Saf. Mobil. Vulnerable Road Usears Pedestr. Bicycl. Motorcycl. 44, 3–11. https://doi.org/10.1016/j.aap.2010.12.007
Zeng, W., Chen, P., Nakamura, H., Iryo-Asano, M., 2014. Application of social force model to pedestrian behavior analysis at signalized crosswalk. Transp. Res. Part C Emerg. Technol. 40, 143–159. https://doi.org/10.1016/j.trc.2014.01.007
Zgonnikov, A., Abbink, D., Markkula, G., 2024. Should I Stay or Should I Go? Cognitive Modeling of Left-Turn Gap Acceptance Decisions in Human Drivers. Hum. Factors 66, 1399–1413. https://doi.org/10.1177/00187208221144561
Zhang, D.C., Highhouse, S., Nye, C.D., 2019. Development and validation of the General Risk Propensity Scale (GRiPS). J. Behav. Decis. Mak. 32, 152–167. https://doi.org/10.1002/bdm.2102
Zhang, W., Luck, S.J., 2008. Discrete fixed-resolution representations in visual working memory. Nature 453, 233–235. https://doi.org/10.1038/nature06860

Zhao, H., Meng, M., Li, X., Xu, J., Li, L., Galland, S., 2024. A survey of autonomous driving frameworks and simulators. Adv. Eng. Inform. 62, 102850. https://doi.org/10.1016/j.aei.2024.102850
Zhao, J., Malenje, J.O., Tang, Y., Han, Y., 2019. Gap acceptance probability model for pedestrians at unsignalized mid-block crosswalks based on logistic regression. Accid. Anal. Prev. 129, 76–83. https://doi.org/10.1016/j.aap.2019.05.012
Zhao, M., Li, R., Xiang, S., Liu, N., 2024. Two different mirror neuron pathways for social and non-social actions? A meta-analysis of fMRI studies. Soc. Cogn. Affect. Neurosci. 19, nsae068. https://doi.org/10.1093/scan/nsae068
Zhi-Xuan, T., Mann, J.L., Silver, T., Tenenbaum, J.B., Mansinghka, V.K., 2020. Online Bayesian goal inference for boundedly-rational planning agents, in: Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20. Curran Associates Inc., Red Hook, NY, USA, pp. 19238–19250.
Zhuang, X., Wu, C., 2011. Pedestrians' crossing behaviors and safety at unmarked roadway in China. Accid. Anal. Prev. 43, 1927–1936. https://doi.org/10.1016/j.aap.2011.05.005

Appendix A

Table A.1 - 4-way factorial ANOVA

Source	Sum Sq.	d.f.	Mean Sq.	F	p	Partial Eta-Sq.
PID	81674	29	2816	2.37	0.0094	0.211
Segment	45562	3	15187	24.54	<0.0001	0.130
Initial Speed	132145	2	66073	132.68	<0.0001	0.302
Initial TTA	57105	1	57105	98.67	<0.0001	0.158
Manoeuvre	1093667	2	546833	589.59	<0.0001	0.782
PID: Segment	53841	87	618	0.9	0.705	0.150
PID: Initial Speed	28882	58	498	1.09	0.375	0.086
PID: Initial TTA	16784	29	578	1.17	0.3166	0.052
PID: Manoeuvre	53793	58	927	1.31	0.1165	0.150
Segment: Initial Speed	43430	6	7238	21.37	<0.0001	0.125
Segment: Initial TTA	17362	3	5787	15.73	<0.0001	0.054
Segment: Manoeuvre	614618	6	102436	202.92	<0.0001	0.668
Initial Speed: Initial TTA	1864	2	932	3.28	0.0448	0.006
Initial Speed: Manoeuvre	8852	4	2213	6.21	0.0001	0.028
Initial TTA: Manoeuvre	5518	2	2759	7.55	0.0012	0.018
PID: Segment: Initial Speed	58948	174	338	1.3	0.0084	0.162
PID: Segment: Initial TTA	32008	87	367	1.41	0.0093	0.095
PID: Segment: Manoeuvre	87838	174	504	1.94	<0.0001	0.223
PID: Initial Speed: Initial TTA	16493	58	284	1.09	0.2997	0.051
PID: Initial Speed: Manoeuvre	41365	116	356	1.37	0.0078	0.119
PID: Initial TTA: Manoeuvre	21191	58	365	1.4	0.0268	0.065
Segment: Initial Speed: Initial TTA	3509	6	584	2.25	0.0368	0.011
Segment: Initial Speed: Manoeuvre	12131	12	1010	3.88	<0.0001	0.038

Segment: Initial TTA: Manoeuvre	3598	6	599	2.3	0.0324	0.012
Initial Speed: Initial TTA: Manoeuvre	2692	4	673	2.58	0.0356	0.009
Error	305222	1172	260			
Total	2840104	2159				

Table A.2 - Pairwise comparisons of progressive time segments in constant speed manoeuvres

Group A	Group B	Lower Limit	A-B	Upper Limit	p
Seg1, Constant	*Seg2, Constant*	0.992	6.736	12.480	0.005
Seg2, Constant	*Seg3, Constant*	13.214	18.958	24.702	<0.001
Seg3, Constant	*Seg4, Constant*	16.721	22.465	28.209	<0.001

Table A.3 - Pairwise comparisons of progressive time segments in deceleration manoeuvres

Group A	Group B	Lower Limit	A-B	Upper Limit	p
Seg1, Deceleration	*Seg2, Deceleration*	-14.285	-8.541	-2.797	<0.001
Seg2, Deceleration	*Seg3, Deceleration*	-17.653	-11.909	-6.165	<0.001
Seg3, Deceleration	*Seg4, Deceleration*	-14.077	-8.333	-2.589	<0.001

Table A.4 - Pairwise comparisons of progressive time segments in mixed manoeuvres

Group A	Group B	Lower Limit	A-B	Upper Limit	p
Seg1, Mixed	*Seg2, Mixed*	-18.035	-12.291	-6.547	<0.001
Seg2, Mixed	*Seg3, Mixed*	-27.063	-21.319	-15.575	<0.001
Seg3, Mixed	*Seg4, Mixed*	-27.480	-21.736	-15.992	<0.001

Table A.5 - Pairwise comparisons of driving manoeuvres at the same segment levels

Group A	Group B	Lower Limit	A-B	Upper Limit	p-value
Seg1, Constant	*Seg1, Deceleration*	-25.327	-19.583	-13.839	<0.001
Seg2, Constant	*Seg2, Deceleration*	-40.605	-34.861	-29.117	<0.001
Seg3, Constant	*Seg3, Deceleration*	-71.473	-65.729	-59.985	<0.001
Seg4, Constant	*Seg4, Deceleration*	-102.271	-96.527	-90.783	<0.001
Seg1, Constant	*Seg1, Mixed*	3.839	9.583	15.327	<0.001
Seg2, Constant	*Seg2, Mixed*	-15.188	-9.444	-3.700	<0.001
Seg3, Constant	*Seg3, Mixed*	-55.466	-49.722	-43.978	<0.001
Seg4, Constant	*Seg4, Mixed*	-99.667	-93.923	-88.179	<0.001
Seg1, Deceleration	*Seg1, Mixed*	23.422	29.166	34.910	<0.001
Seg2, Deceleration	*Seg2, Mixed*	19.672	25.416	31.160	<0.001
Seg3, Deceleration	*Seg3, Mixed*	10.262	16.006	21.751	<0.001
Seg4, Deceleration	*Seg4, Mixed*	-3.139	2.604	8.348	1

Table A.6 - Linear mixed-effects model analysis

Model information:							
Number of observations				2400			
Fixed effects coefficients				32			
Random effects coefficients				30			
Covariance parameters				2			
Formula:							
Linear Mixed Formula with 5 predictors							
Model fit statistics:							
AIC		*BIC*		*LogLikelihood*		*Deviance*	
21352		21548		-10642		21284	
Fixed Effects Coefficients (95% Confidence Intervals):							
Name	*Estimate*	*SE*	*tStat*	*DF*	*p*	*Lower*	*Upper*
(Intercept)	78.52	7.92	9.90	2368	<0.001	62.98	94.07
Segment	-2.5	2.89	-0.86	2368	0.388	-8.17	3.17
Speed	-1.31	0.18	-7.28	2368	0	-1.66	-0.96
TTA	5.11	1.67	3.06	2368	0.002	1.83	8.39
Constant speed	-20.16	13.56	-1.48	2368	0.137	-46.76	6.44
Deceleration	28.62	13.56	2.11	2368	0.035	2.01	55.23
Two-stage deceleration	-55.19	13.88	-3.97	2368	0	-82.42	-27.97
Segment: Speed	0.24	0.06	3.64	2368	0	0.11	0.37
Segment: TTA	-1.36	0.61	-2.23	2368	0.025	-2.56	-0.17
Speed: TTA	0.03	0.03	0.87	2368	0.382	-0.04	0.11
Segment: Constant speed	-13.07	4.95	-2.63	2368	0.008	-22.79	-3.36
Segment: Deceleration	2.49	4.95	0.50	2368	0.615	-7.22	12.21
Segment: Two-stage deceleration	26.79	5.07	5.28	2368	0	16.85	36.73
Speed: Constant speed	0.14	0.31	0.46	2368	0.643	-0.46	0.76
Speed: Deceleration	-0.75	0.31	-2.43	2368	0.015	-1.37	-0.15
Speed: Two-stage deceleration	0.38	0.31	1.22	2368	0.222	-0.23	0.99
TTA: Constant speed	8.02	2.86	2.80	2368	0.005	2.41	13.63
TTA: Deceleration	-8.07	2.86	-2.82	2368	0.005	-13.68	-2.47
TTA: Two-stage deceleration	4.74	2.92	1.62	2368	0.105	-0.99	10.49

Segment: Speed: TTA	-0.002	0.01	-0.17	2368	0.859	-0.03	0.02
Segment: Speed: Constant speed	0.06	0.11	0.55	2368	0.578	-0.16	0.28
Segment: Speed: Deceleration	0.21	0.11	1.86	2368	0.062	-0.01	0.43
Segment: Speed: Two-stage deceleration	-0.14	0.11	-1.27	2368	0.204	-0.36	0.08
Segment: TTA: Constant speed	-0.96	1.04	-0.91	2368	0.358	-3.00	1.08
Segment: TTA: Deceleration	1.68	1.04	1.60	2368	0.108	-0.36	3.72
Segment: TTA: Two-stage deceleration	-1.89	1.06	-1.77	2368	0.077	-3.98	0.2
Speed: TTA: Constant speed	-0.08	0.066	-1.23	2368	0.219	-0.21	0.04
Speed: TTA: Deceleration	0.16	0.066	2.5	2368	0.012	0.03	0.29
Speed: TTA: Two-stage deceleration	-0.11	0.066	-1.74	2368	0.081	-0.24	0.01
Segment: Speed: TTA: Constant speed	-0.002	0.024	-0.07	2368	0.944	-0.04	0.04
Segment: Speed: TTA: Deceleration	-0.03	0.024	-1.42	2368	0.153	-0.08	0.01
Segment: Speed: TTA: Two-stage deceleration	0.04	0.024	1.69	2368	0.091	-0.006	0.08

Random Effects Covariance Parameters (95% Confidence Intervals):

Group: PID (30 Levels)

Name1	*Name2*	*Type*	*Estimate*	*Lower*	*Upper*
(Intercept)	*(Intercept)*	std		NaN	NaN

Group: Error

Name	*Estimate*	*Lower*	*Upper*
Res Std	20.392	19.823	20.977

Appendix B

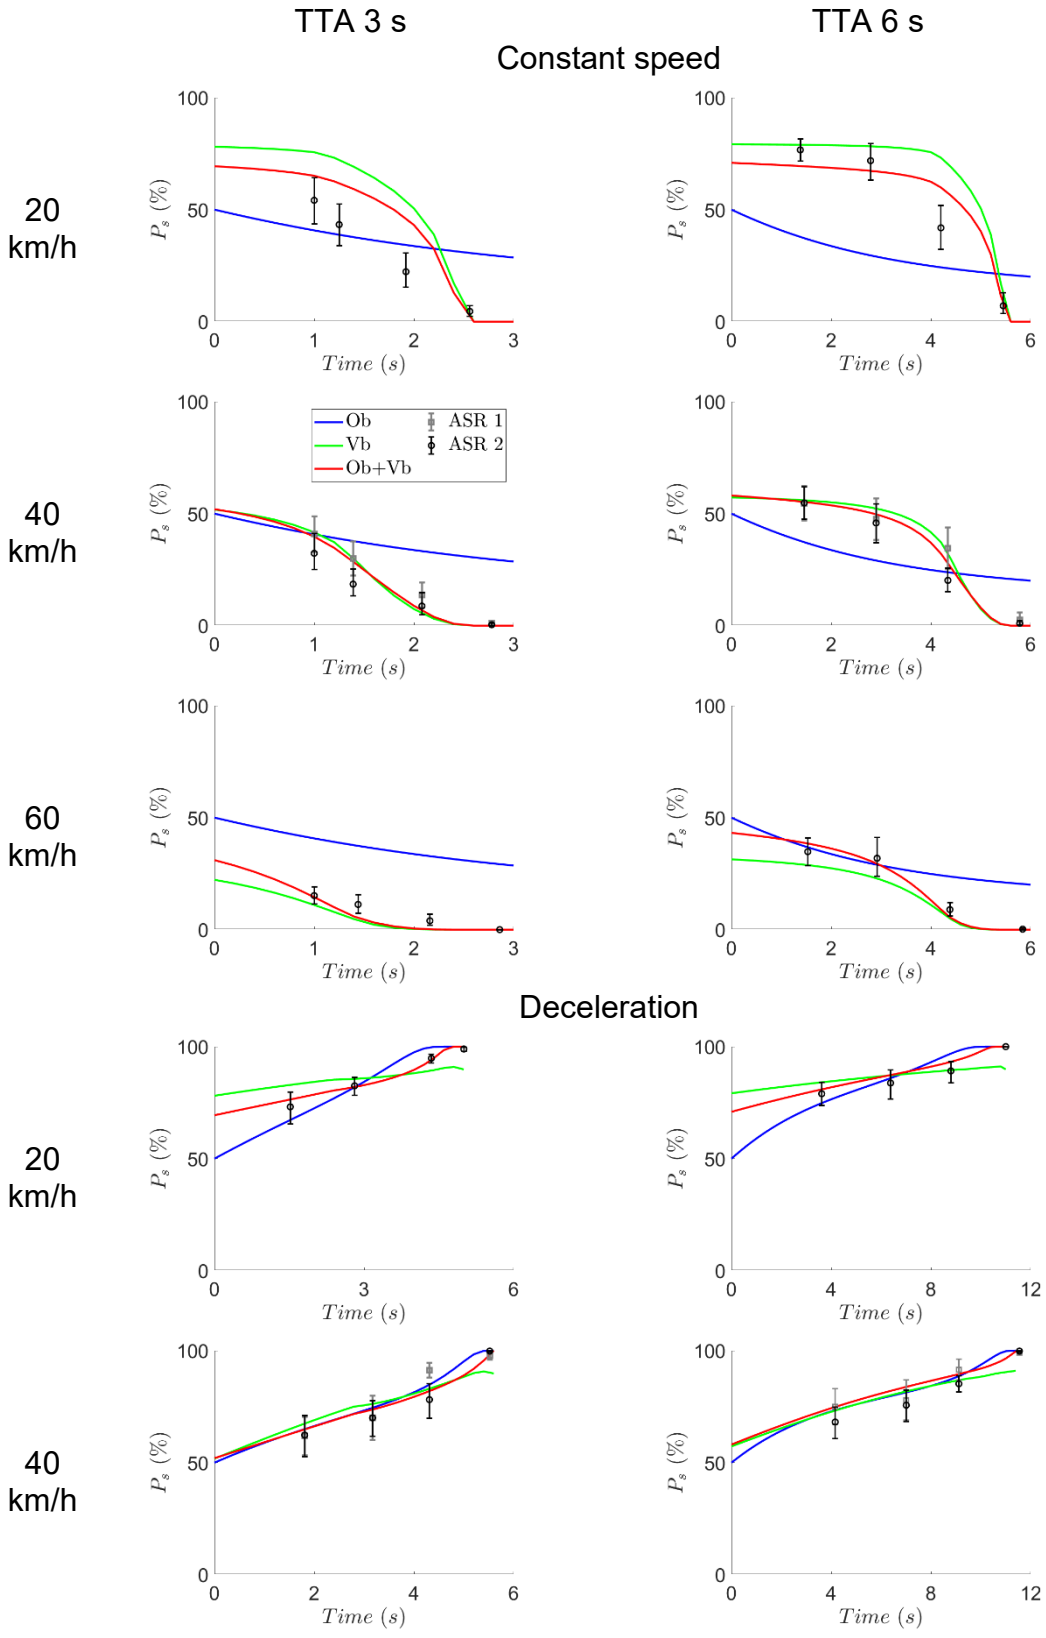
TTA 3 s
TTA 6 s
Constant speed
20 km/h
40 km/h
60 km/h
Deceleration
20 km/h
40 km/h
Ob
Vb
Ob+Vb
ASR 1
ASR 2
P_s (%)
Time (s)

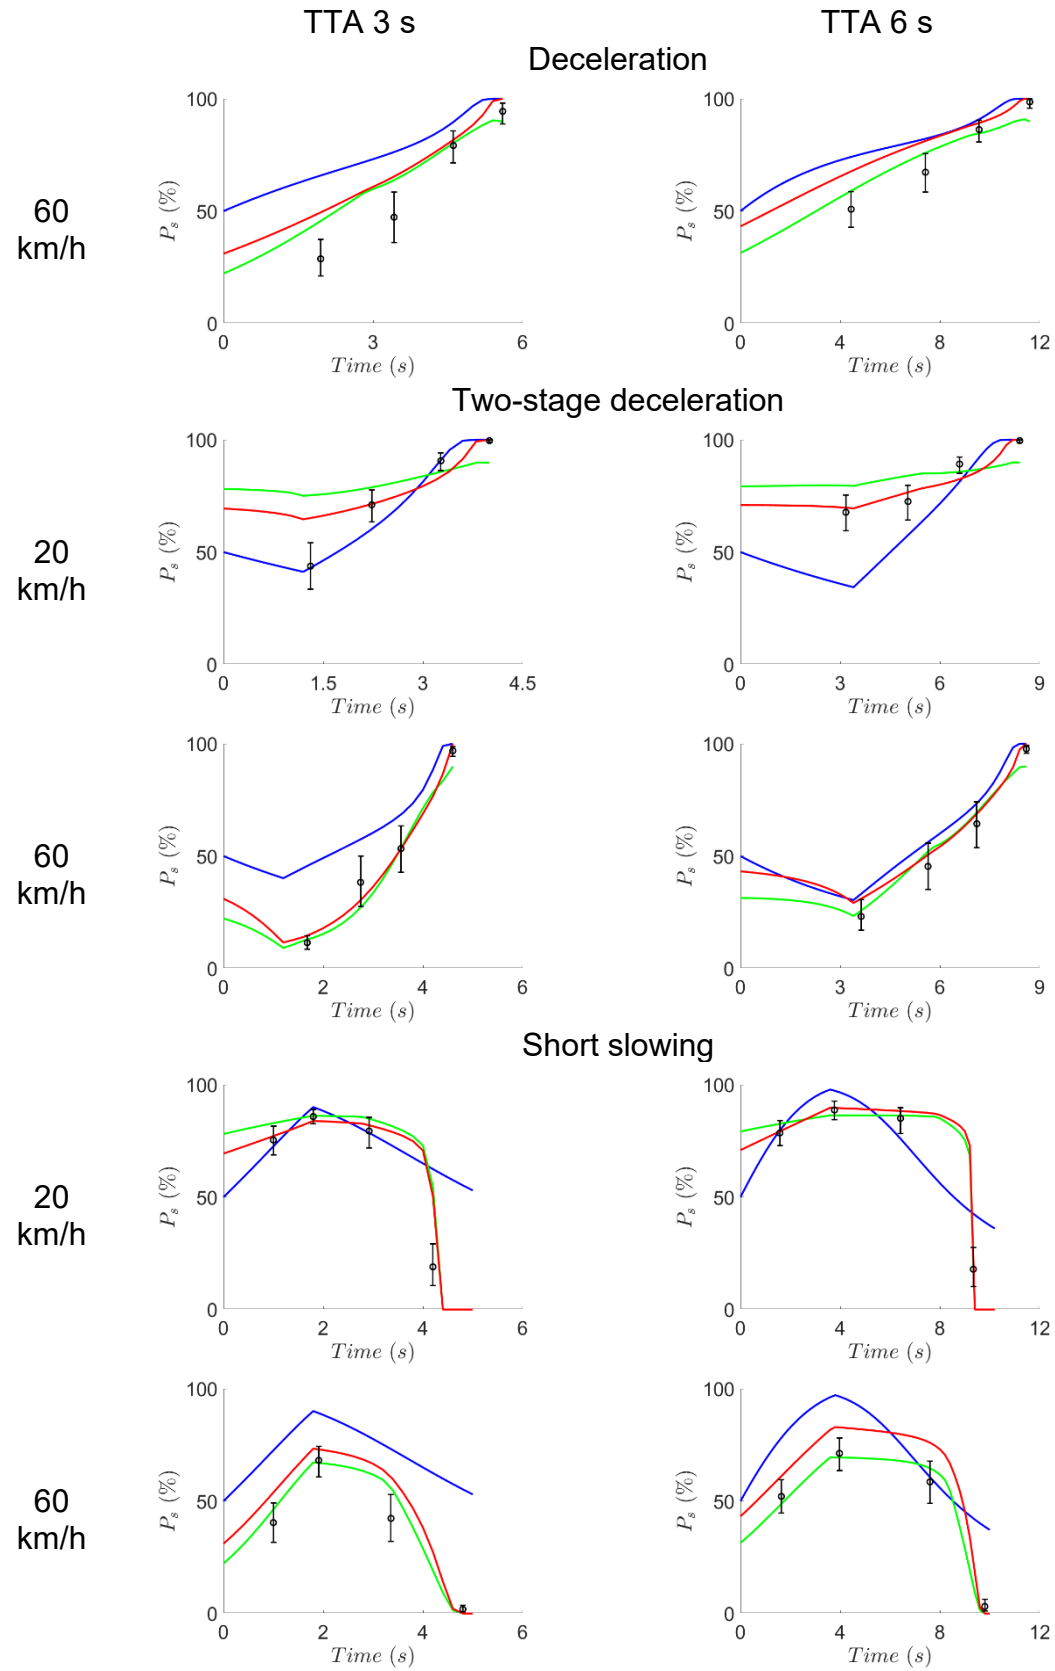

Figure B.1 - Comparison between model predictions and average pedestrian beliefs for all 20 kinematic scenarios (parameter settings of Chapter 3)

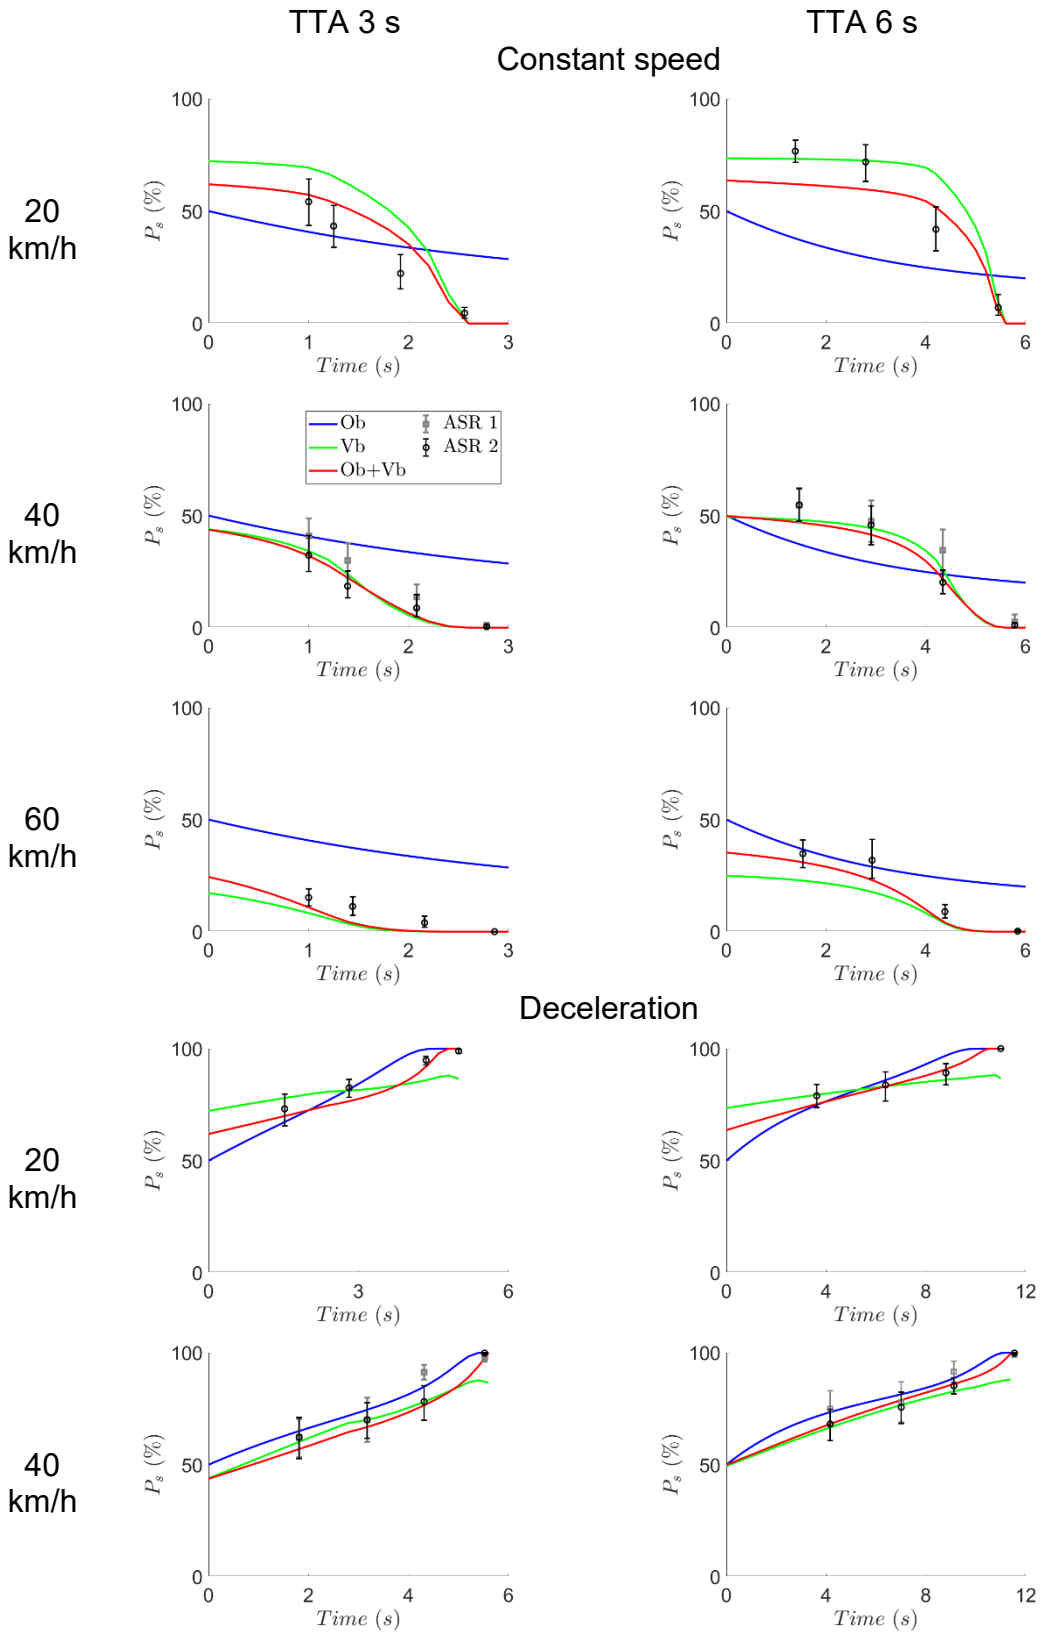
TTA 3 s
TTA 6 s
Constant speed
20 km/h
40 km/h
60 km/h
Deceleration
20 km/h
40 km/h
Ob
Vb
Ob+Vb
ASR 1
ASR 2
P_s (%)
Time (s)

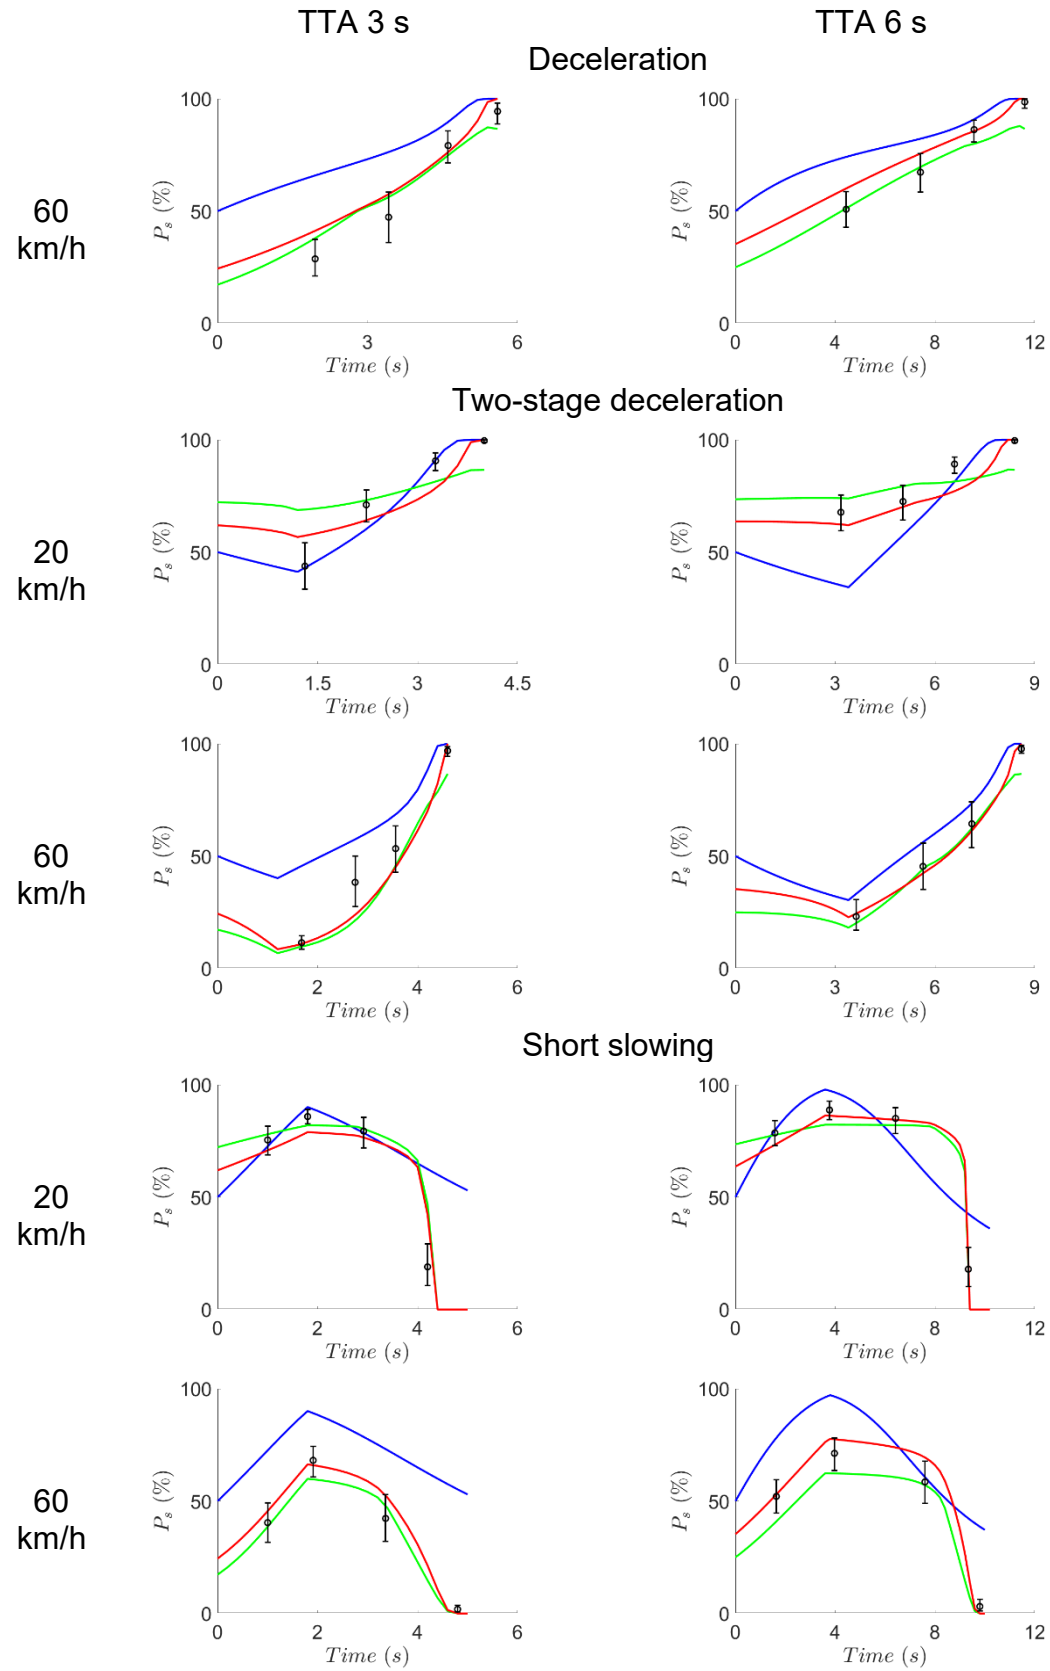

Figure B.2 - Comparison between model predictions and average pedestrian beliefs for all 20 kinematic scenarios (refitted only parameter M)

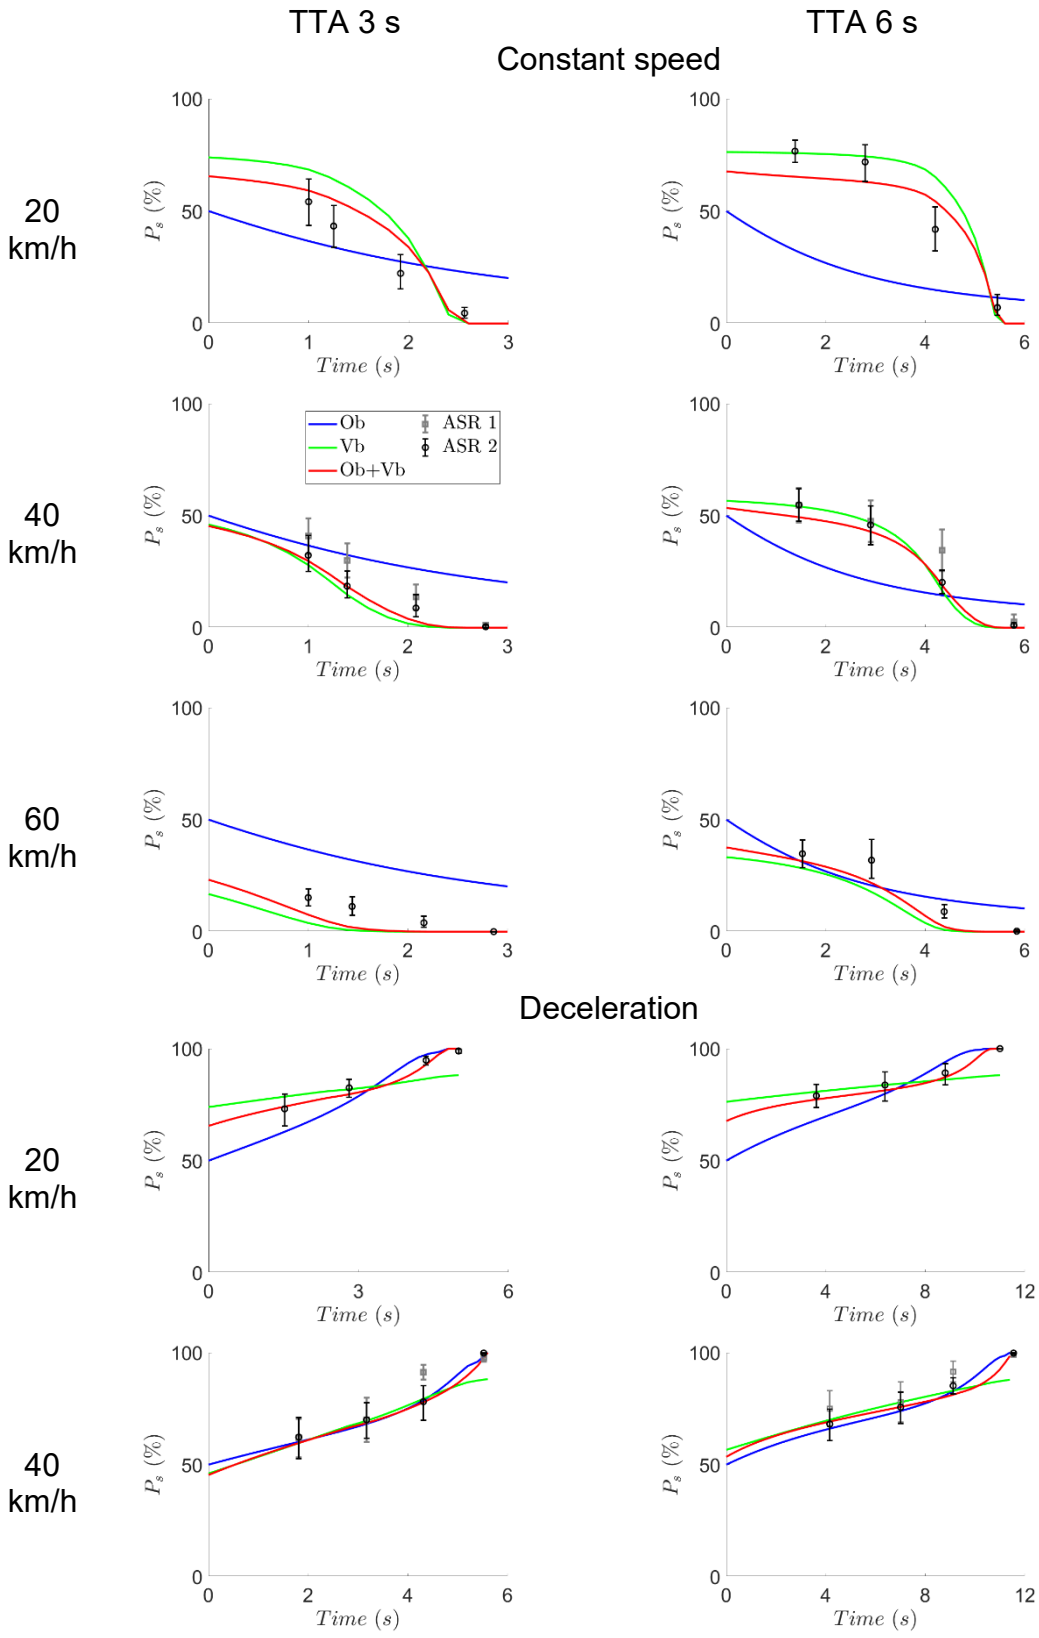
TTA 3 s
TTA 6 s
Constant speed
20 km/h
40 km/h
60 km/h
Deceleration
20 km/h
40 km/h
Ob
Vb
Ob+Vb
ASR 1
ASR 2
P_s (%)
Time (s)

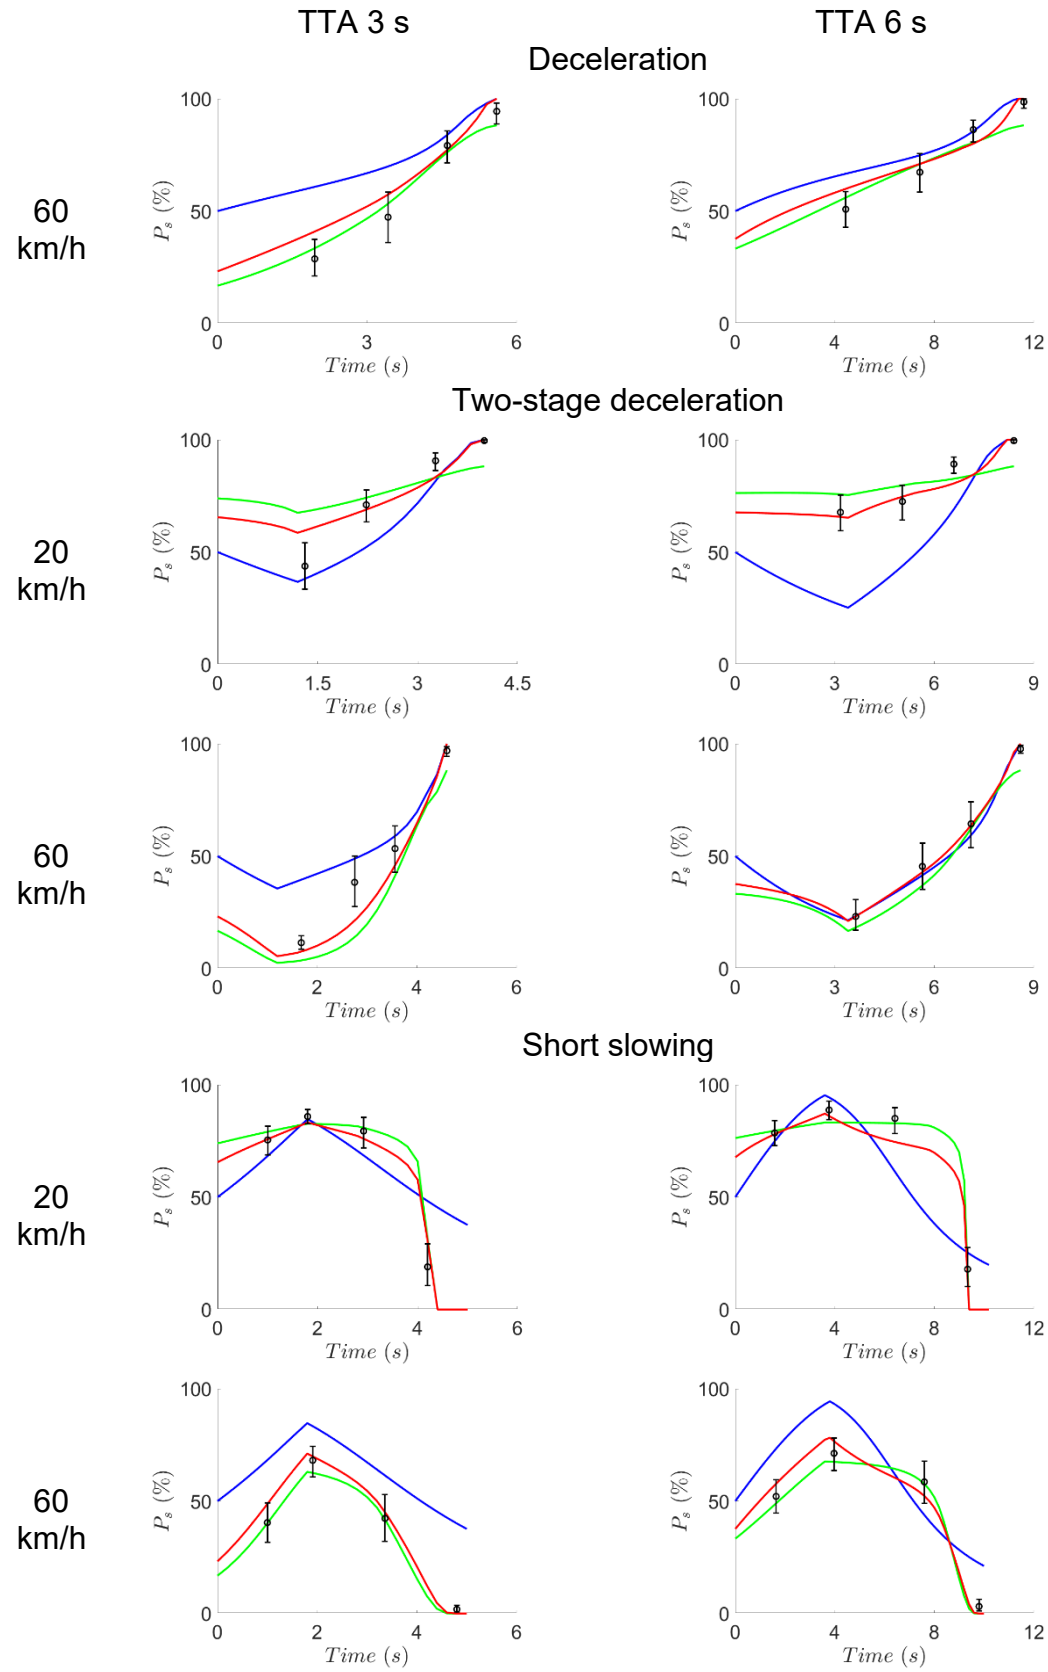

Figure B.3 - Comparison between model predictions and average pedestrian beliefs for all 20 kinematic scenarios (refitted all parameters)

Appendix C

Behaviour Estimation Task Post-Trial Questions:

1. Was the vehicle stopping for you or was it maintaining its speed and passing you?

Stopping	**Passing**

2. How confident are you in your previous answer? Please rate your confidence level on a scale from 1 to 9.

1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
Not confident at all				Somewhat confident				Totally confident

Appendix D

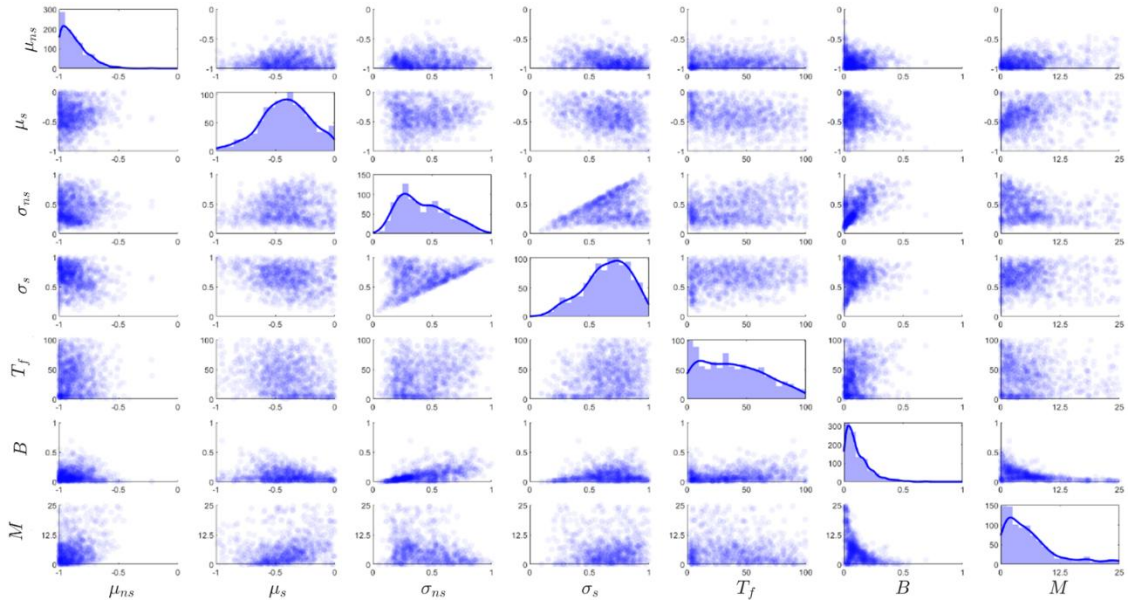

Figure D.1 - Ob's pairwise parameter scatterplot matrix. The histograms in the diagonal illustrate the distribution of the respective parameter values with the BSCV-obtained RMSEs. The scatterplots show the pairwise parameter combination areas with obtained RMSEs

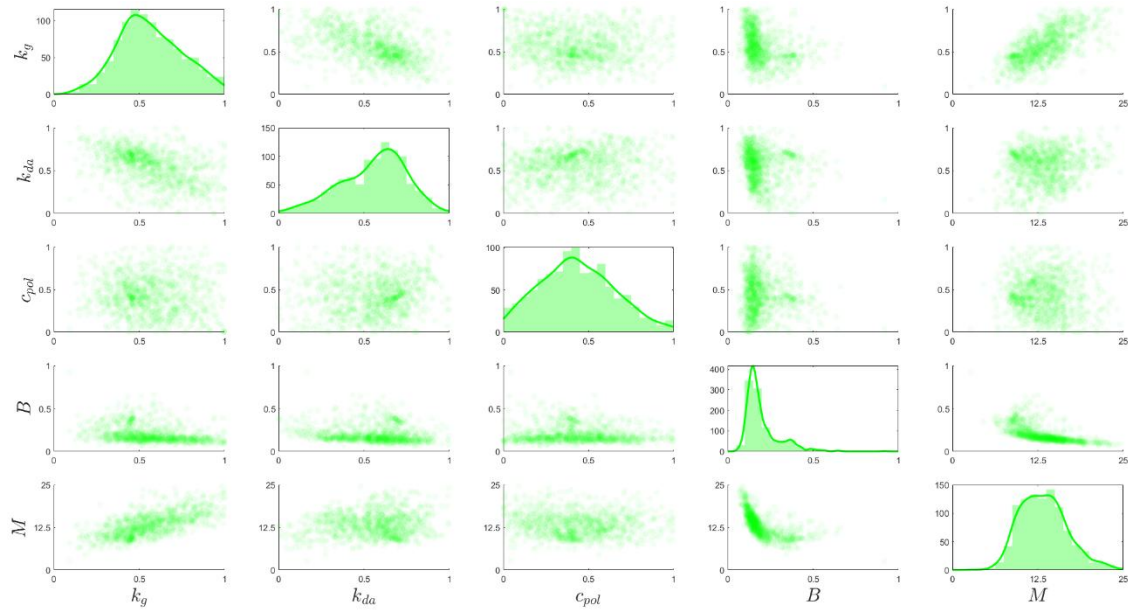

Figure D.2 - Vb's pairwise parameter scatterplot matrix. The histograms in the diagonal illustrate the distribution of the respective parameter values with the BSCV-obtained RMSEs. The scatterplots show the pairwise parameter combination areas with obtained RMSEs

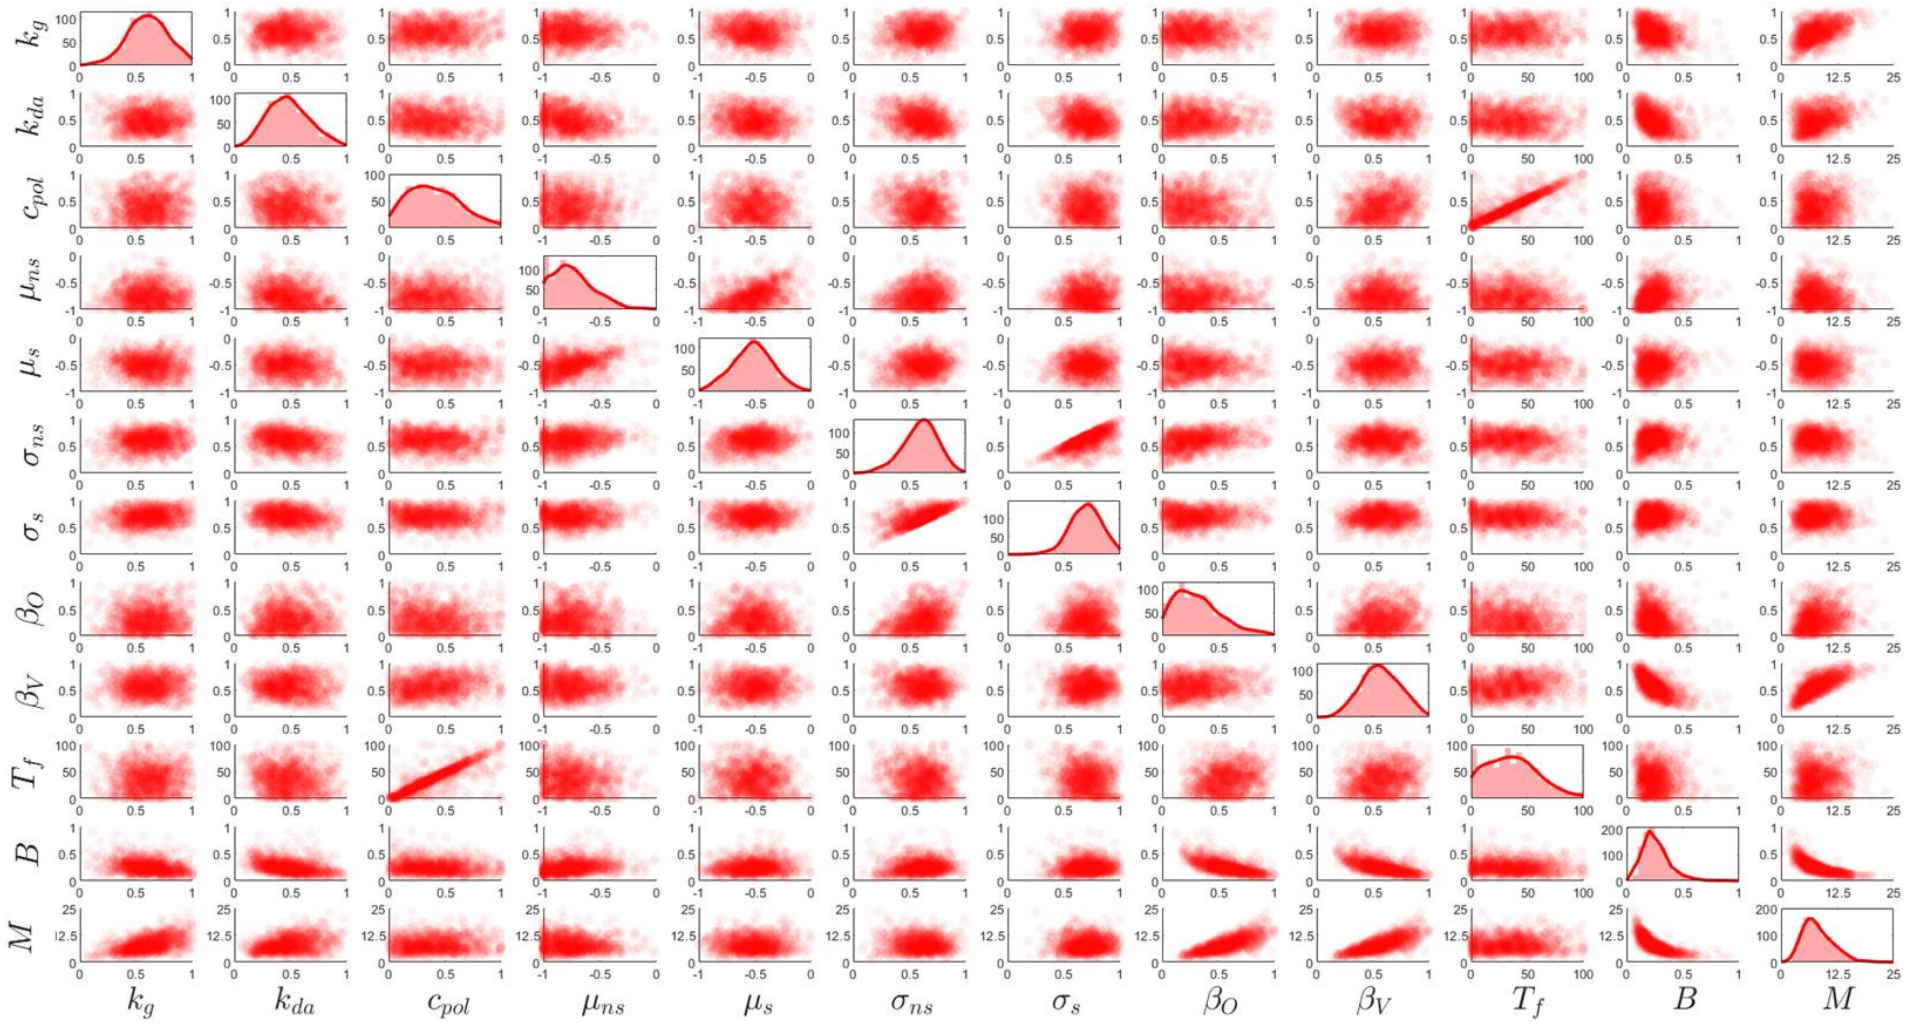

Figure D.3 - Ob+Vb's pairwise parameter scatterplot matrix. The histograms in the diagonal illustrate the distribution of the respective parameter values with the BSCV-obtained RMSEs. The scatterplots show the pairwise parameter combination areas with obtained RMSEs

Appendix E

Table E.1 - Models of pedestrian decision-making

Study	Model	Input	Output
Kotseruba and Rasouli (2023)	Agent-based model: Intend-Wait-Perceive-Cross	Vehicle kinematics (speed, position, acceleration), Pedestrian characteristics, Perception parameters (field of view - FoV), Scanning strategy, Time since last observation	Pedestrian crossing decisions (e.g., Cross or Wait), Waiting time, Number of collisions, Head turns, Minimum TTC
Sun et al. (2002)	Binary Logit Model	Gap size, waiting time, Pedestrian age, Pedestrian gender, Number of pedestrians waiting	Probability of Gap Acceptance (Binary decision: accept or reject)
Zhao et al. (2019)	Binary Logit Model	Gap size, Crossing distance, Waiting time, Position of pedestrian in relation to the kerb	Probability of Gap Acceptance (Binary decision: accept or reject)
Himanen and Kulmala (1988)	Multinomial Logit (MNL) Model	Pedestrian distance from kerb/refuge, Locality (City size), Pedestrian group size, Vehicle speed, Number of vehicles approaching, Vehicle type, Pedestrian on marked crossing, Driving direction, Sex of pedestrian	Probability of continuing (walks on) or reacting (stops, retreats, or needs to run)

Tian et al. (2022)	Psychophysics-based Gap Acceptance (PGA) Model with mixed effects (Logit Model based on Visual Looming cue).	PGA Model Input: $\dot{\theta}$, Looming calculation inputs: TTA, Vehicle speed, Vehicle length, Vehicle width, Lateral distance from car to pedestrian, Distance between pedestrian and vehicle	Probability of Gap Acceptance, Post-encroachment time, Crossing duration, and Percentage of unsafe decisions and tight fits
Raghuram Kadali et al. (2014)	Artificial Neural Network (ANN) Model	Gap size, Frequency of attempt, Rolling gap (Yes/No), Speed changes condition (Yes/No), Vehicle speed, Movement of pedestrian, Group size, Waiting time, Near or far lane gap, Type of vehicle	Probability of Gap Acceptance (Binary decision: accept or reject)
Pawar and Patil (2016)	Binary Logit Model	Gap Time or Distance, Speed of vehicle, Traffic volume	Probability of Gap Acceptance (Binary decision: accept or reject)
Tian et al. (2025)	PT-PRD Model: Hybrid Perception Strategy and Crossing Initiation Model	Visual Perception Cues: $\dot{\theta}$ and $\dot{\tau}$, Physical Kinematic Inputs: Vehicle distance, Vehicle width, Vehicle speed, Vehicle deceleration, Time gap, TTC	Binary Crossing Decision, Crossing Initiation Time, Cumulative Crossing Probabilities

Wang et al. (2023)	Deep Q-Networks Learning, Constrained by Bayesian Perceptual Filtering of Noisy Visual Input	Position and Velocity of vehicle and pedestrian, Position and Velocity uncertainty, Visual noise	Discrete decision (Go or Not Go), Gap Acceptance Rate, Crossing Initiation Time
Crosato et al. (2023)	Deep Reinforcement Learning Framework using Social Value Orientation Informed Reward Function, Social Force Model and Gap Acceptance Model, incorporating Situational Awareness	Reward Function Input: AV performance parameters, Pedestrian intentions/comfort, Pedestrian crossing speed, Motivation, Distance, Social Value Orientation, Pedestrian Model Inputs: AV velocity, Distance, Lane width, AV acceleration, Pedestrian desired walking speed	AV Decision-Making Policy (Driving Styles: Egoistic to Pro-Social), AV Longitudinal acceleration/deceleration, Pedestrian Trajectory
Markkula et al. (2023)	Integrated Computational Psychological Model (Bayesian perception, Theory of mind/Behaviour estimation, Affordance-based long-term value estimation, and Evidence accumulation decision-making)	Noisy Position and Speed of other agent, Acceleration/Speed adjustments	Continuous Action Decisions

Giles et al. (2019)	Variable-Drift Diffusion Model (VDDM)	Apparent TTA (τ) and TTA rate of change ($\dot{\tau}$)	Crossing Initiation Time
Pekkanen et al. (2022)	Variable-Drift Diffusion Model (VDDM)	TTA (τ), Distance, TTA rate of change ($\dot{\tau}$), Presence of eHMI (Yes/No), Noise	Decision Timing, Crossing Onset Time Distribution
Markkula et al. (2018)	Evidence Accumulation/DDM)	Looming, TTA (τ) and TTA rate of change ($\dot{\tau}$), Presence of explicit signal (Yes/No), Vehicle's Awareness of pedestrian	Pedestrian Crossing/Yielding Decisions