Assessment of Latent Pedestrian–Vehicle Interaction Risk Profiles at Midblock Crossing in VR

Rulla Al-Haideri*[1], Bilal Farooq[2], and Elisabetta Cherchi[3]

[1]NSERC Postdoctoral Fellow, Laboratory of Innovations in Transportation (LiTrans), Toronto Metropolitan University, Canada
[2] Associate Professor, Canada Research Chair in Disruptive Transportation Technologies and Services, Laboratory of Innovations in Transportation (LiTrans), Toronto Metropolitan University, Canada
[3] Professor, Department of Civil and Urban Engineering, New York University Abu Dhabi, UAE

Short summary

Pedestrian safety at midblock crossings is a critical concern in mixed traffic environments where autonomous vehicles (AVs) and human-driven vehicles (HDVs) share the road. Pedestrians often infer intent from vehicle motion in AV encounters, making them vulnerable to small shifts in conflict margins. This study investigates whether virtual reality (VR) crossing sessions separate into distinct interaction risk profiles and whether AV-only sessions shift profile prevalence compared to HDV-only sessions. Using large-scale immersive VR experiments from Toronto, Canada, and Newcastle, England, we compute surrogate safety measures (SSMs) and apply latent profile analysis (LPA) to identify distinct pedestrian crossing stances, ranging from risk-accepting to highly cautious. Key findings show that Newcastle exhibits a higher prevalence of high-urgency risk profiles in AV-only sessions, indicating that AVs contribute to higher-risk encounters. In contrast, Toronto shows no significant difference between AV-only and HDV-only sessions, suggesting that contextual factors influence the impact of AVs on pedestrian safety.
Keywords: surrogate safety measures, latent profile analysis, virtual reality, autonomous vehicles, pedestrian crossings.

1 Introduction

A traffic conflict is an observable interaction in which a collision would have happened if at least one road user had not taken an evasive action (*e.g.*, braking, swerving, stopping, or accelerating) Laureshyn & Várhelyi (2018). These conflicts vary in severity, depending on the time and space margins between road users. When these margins become small enough, crash avoidance may depend on rapid or substantial evasive manoeuvres. In this case, the interaction becomes safety-critical and could escalate to a crash if avoidance is no longer possible. This severity is commonly operationalized using surrogate safety measures (SSMs). SSMs occur more frequently than crashes and preserve the kinematic structure of interactions. This makes them suitable for comparing safety-relevant behaviour under different conditions. A key challenge in pedestrian–vehicle safety evaluation is heterogeneity. A small share of encounters can dominate the safety-relevant tail. Therefore collapsing SSMs to a single mean or median can mask distinct risk regimes including rare but consequential near misses. This issue becomes more prominent in autonomous vehicles (AVs) settings. In such encounters, pedestrians often cannot rely on driver eye contact or informal gestures and must infer intent primarily from vehicle motion and any external human–machine interface (eHMI) cues. Immersive virtual reality (VR) proides a safe and a controlled way to study pedestrian–vehicle interactions while preserving naturalistic movement and decision-making which support SSM-based safety assessment. Prior work employed latent constructs and segmentation to capture heterogeneity in pedestrian behaviour and safety outcomes Hosseini Shoabjareh et al. (2021). Recently, LPA has been used to stratify interaction severity when multiple SSMs are considered jointly Maji & Ghosh (2025). However, midblock pedestrian–vehicle interactions particularly in pedestrian–AV contexts are still evaluated using single summary statistic. This aggregation can mask whether AV-only versus human-driven vehicle (HDV)-only scenarios shift the distribution of interaction risk, especially in the upper tail where safety margins are smallest.

This paper employs LPA as an exploratory segmentation method to summarize multi-indicator SSM heterogeneity to a small number of interpretable risk profiles. The use of LPA is justified using two points. First, LPA moves the evaluation beyond single indicator averages by identifying distinct urgency regimes in the joint SSM space. Second, it enables direct comparisons of how often high urgency regimes occurred across vehicle types and locations. We apply this approach to a large scale VR midblock crossing experiments conducted in Toronto, Canada and Newcastle, England. Our unit of analysis is the VR experiment session with unique variable values (*e.g.*, low vehicle flow, snowy weather, and a median in the middle), summarized using session level conflict indicators. For details on experiments, sessions, and data collection campaign, we refer the reader to Nazemi et al. Nazemi et al. (2025). We estimated LPA separately by location using two temporal indicator specifications: Model A used min(MTTC) and the corresponding CTTC at the same time step, while Model B used min(T_2) and the corresponding CTTC. We then compared profile prevalence between AV-only and HDV-only sessions within each location by focusing on the top highest risk end of the profile structure. The main objective of this paper is to identify and compare latent interaction risk profiles in VR midblock pedestrian–vehicle encounters using SSMs, and to examine how profile prevalence and structure vary by vehicle type and location. This paper addresses the following research questions:

1. Within each location, what latent interaction risk profiles emerge when each VR session is represented by multiple temporal SSMs?
2. Within each location, does the proportion of sessions assigned to the highest risk profile differ between AV-only and HDV-only scenarios?
3. Across locations, are the latent profiles comparable in number and indicator patterns, and do AV-only versus HDV-only differences change when the indicator set is used?

The remainder of the paper is organized as follows. Section 2 presents the methods, including the data analysis and modelling framework. Section 3 reports the results and discussion. Section 4 concludes and outlines future work.

2 Methods

Data Analysis

This study uses two VR datasets collected in Toronto (109 participants) and Newcastle (249 participants). Each dataset includes time-stamped trajectories for pedestrians and simulated vehicle agents during a midblock crossing task on a two-lane road. Participants complete repeated sessions and each lasts for 60 s under systematically varied conditions. Within the two locations, each participant completed 12 to 24 sessions, with conditions varying by time of day, weather, vehicle type (AV versus HDV), presence of an eHMI for the AV, presence of a median, behaviour of avatar pedestrians, and traffic flow. The scenario structure is comparable across locations to support within location modelling under consistent task demands. Three predictive temporal conflict indicators are calculated to quantify session-level interaction criticality under both collision course and non-collision course encounters. These include the modified time-to-collision (MTTC), time-to-arrival (T_2), and closing time-to-collision (CTTC). For LPA, each 60 s session is represented by its most critical time step. This is defined as the minimum of a primary temporal indicator and the corresponding CTTC at that same time step. Two indicator sets are analyzed within each location. Model A uses min(MTTC) with CTTC evaluated at the time of min(MTTC). Model B uses min(T_2) with CTTC evaluated at the time of min(T_2). MTTC is the time for two road users to collide under constant acceleration extrapolation, evaluated only when the pair is on a collision course Ozbay et al. (2008). In this paper, road users are represented as rectangles and collision times are computed using the corner–side procedure of Laureshyn et al. Laureshyn et al. (2010). T_2 is the predicted arrival time of the later road user to the avoided conflict point Laureshyn et al. (2010). CTTC is the instantaneous closing time implied by the current separation distance divided by the line-of-sight closing speed (relative velocity projected onto the line of sight) under constant-velocity extrapolation Al-Haideri et al. (2025). The estimation samples used in this study represent subsets of the full recorded datasets. Sessions are included in the LPA estimation sample only when both indicators required by the selected model are defined (finite after pre-processing).

Modelling Framework

The VR procedure and control variables are coded consistently across Toronto and Newcastle. However, we estimate LPA separately by location because pooling requires assuming that the indicator distributions defining each profile are the same across sites. Testing and imposing cross site measurement invariance within profiles (and estimating a pooled multi-group LPA) is beyond the scope of this paper and is left for future work. In each location, models with $K = 1, 2, \ldots, K_{\max}$ are evaluated, and solutions are discarded if any profile contained fewer than 5% of sessions. Classification quality is assessed using posterior membership probabilities. The distribution of maximum posterior probabilities is summarized, and average posterior probabilities of assignment of around 0.70 or higher are treated as indicative of acceptable classification Nagin & Odgers (2010). Among candidate solutions, the smallest number of profiles K that still provides an adequate fit is selected. First, the model with the lowest BIC is identified. The smallest K whose BIC is within 10 points of that best (lowest) BIC is then selected. The K is increased only if moving from $K-1$ to K improves BIC by at least 2 points. Otherwise the more parsimonious $K-1$ solution is kept. After selecting K, profile prevalence is reported separately for AV-only and HDV-only sessions for each location. Profiles are ranked only for consistent labelling (rank 1 as highest risk). These ranks are not used in model estimation or model selection. To rank profiles, a composite severity score is defined for each session as:

$$r_i = -\sum_{j \in \mathcal{J}} z_{ij}, \tag{1}$$

where $\mathcal{J}$ is the set of standardized indicators used in the selected model. For each session, the LPA model returns posterior membership probabilities for all K profiles. Each session is assigned to the profile with the largest posterior probability. Using these assignments, profile prevalence is computed separately for AV-only and HDV-only sessions per location. For each session i in a given location and indicator model, the 2D feature vector is defined as:

$$\mathbf{x}_i = \begin{bmatrix} z_{\text{loc,model}}(\log\, t_i) \\ z_{\text{loc,model}}(\log\, c_i) \end{bmatrix} \in R^2,$$

where t_i is the session level conflict indicator ($t_i = \min(\text{MTTC})$ for Model A or $t_i = \min(\text{T}_2)$ for Model B), and c_i is the corresponding CTTC value evaluated at the time step where t_i occurred. The $z_{\text{loc,model}}(\cdot)$ denotes z-score standardization computed within each location and indicator model using the sessions included in that model. LPA represents the distribution of the session level indicator vector $\mathbf{x}_i$ as a mixture of K latent profiles. Each session i belongs to an unobserved profile $Z_i \in \{1, \ldots, K\}$ with mixing proportions $\pi_k = \Pr(Z_i = k)$, where $\sum_{k=1}^{K} \pi_k = 1$. The mixture density is expressed by:

$$f(\mathbf{x}_i) = \sum_{k=1}^{K} \pi_k \, \phi(\mathbf{x}_i \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)\,, \tag{2}$$

where $\phi(\cdot \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ is the bivariate normal density for profile k with mean $\boldsymbol{\mu}_k$ and covariance $\boldsymbol{\Sigma}_k$. Model parameters $\{\pi_k, \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}_{k=1}^{K}$ are estimated by maximum likelihood using an expectation–maximization algorithm. For each session, posterior membership probabilities are obtained by:

$$\gamma_{ik} = \Pr(Z_i = k \mid \mathbf{x}_i) = \frac{\pi_k \, \phi(\mathbf{x}_i \mid \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)}{\sum_{h=1}^{K} \pi_h \, \phi(\mathbf{x}_i \mid \boldsymbol{\mu}_h, \boldsymbol{\Sigma}_h)}, \tag{3}$$

and the session was assigned to the profile with the largest γ_{ik}. All models are estimated in R using `tidyLPA` van Lissa et al. (2018).

3 Results and Discussion

Table 1 summarizes the selected number of profiles, estimation sample sizes, and classification diagnostics for each model. The selected solutions ranged from $K = 4$ (Newcastle Model A) to $K = 6$ (Newcastle Model B and Toronto Model A). The entropy summarizes how clearly sessions are assigned to a single latent profile, and values closer to 1 indicate clearer separation. The mean maximum posterior probability is the average across sessions of the largest profile membership probability, where values closer to 1 imply higher assignment certainty. As can be seen from the table, classification separation is strong for Newcastle Model A (entropy = 0.9551), Newcastle

Table 1: LPA model fit summary by location.

Location	Model	K	N_{AV}	N_{HDV}	BIC	Entropy	Min. class prop.	Mean max post.
NC	A	4	445	242	616.84	0.9551	0.1645	0.9830
NC	B	6	1141	569	6024.75	0.9465	0.0696	0.9655
TO	A	6	607	584	2403.77	0.9539	0.0705	0.9677
TO	B	5	445	458	3717.29	0.8641	0.0554	0.9128

Notes: TO: Toronto, NC: Newcastle. K is selected number of latent profiles. N_{AV} and N_{HDV} are numbers of sessions used for estimation in AV-only and HDV-only.

Table 2: Key contrasts in interaction risk profile prevalence between AV-only and HDV-only scenarios by location.

Location	Model	Contrast	k_{AV}	N_{AV}	p_{AV}	k_{HDV}	N_{HDV}	p_{HDV}	Δp	95% CI
NC	A	Top-1 risk profile	67	445	0.1506	46	242	0.1901	-0.0395	[−0.1239, 0.0411]*
NC	A	Top-2 risk profiles (rank 1–2)	200	445	0.4494	77	242	0.3182	0.1313	[0.0245, 0.2332]
NC	B	Top-1 risk profile	357	1141	0.3129	132	569	0.2320	0.0809	[0.0182, 0.1412]
NC	B	Top-2 risk profiles (rank 1–2)	489	1141	0.4286	161	569	0.2830	0.1456	[0.0788, 0.2100]
TO	A	Top-1 risk profile	126	607	0.2076	119	584	0.2038	0.0038	[−0.0611, 0.0685]*
TO	A	Top-2 risk profiles (rank 1–2)	175	607	0.2883	154	584	0.2637	0.0246	[−0.0472, 0.0960]*
TO	B	Top-1 risk profile	29	445	0.0652	21	458	0.0459	0.0193	[−0.0233, 0.0618]*
TO	B	Top-2 risk profiles (rank 1–2)	103	445	0.2315	100	458	0.2183	0.0131	[−0.0638, 0.0899]*

Notes: k_{AV} and k_{HDV} are numbers of sessions assigned to contrasted profile set (top-1 or top-2) in AV-only and HDV-only; N_{AV} and N_{HDV} are total numbers of sessions scenarios. $p_{\mathrm{AV}} = k_{\mathrm{AV}}/N_{\mathrm{AV}}$ and $p_{\mathrm{HDV}} = k_{\mathrm{HDV}}/N_{\mathrm{HDV}}$. $\Delta p = p_{\mathrm{AV}} - p_{\mathrm{HDV}}$ is reported with Newcombe–Wilson 95% CIs. *CI includes 0 (contrast is not distinguishable at the 95% level).

Model B (entropy = 0.9465), and Toronto Model A (entropy = 0.9539), with mean maximum posterior probabilities above 0.96. Toronto Model B shows weaker separation (entropy = 0.8641; mean maximum posterior probability = 0.9128). This suggests that the T_2-based indicator space in Toronto captures either a more continuous spectrum of risk or higher within-profile variability. As a result, profiles overlap more, and session-to-profile assignments are less stable.

AV-only versus HDV-only differences in high criticality profile prevalence

Table 2 presents contrasts in profile prevalence between AV-only and HDV-only sessions for (i) the highest interaction risk profile and (ii) the combined set of the top two highest interaction risk profiles. These contrasts focus on the upper tail of interaction risk, where smaller temporal margins are most safety-relevant. As a general observation, Table 1 indicates strong profile separation in Newcastle and in Toronto under Model A. In contrast, Toronto under Model B shows weaker separation, with lower entropy and a lower mean maximum posterior probability, consistent with greater overlap in the T_2-based indicator space. In Newcastle, AV-only sessions exhibit a higher prevalence of the combined top two highest risk profiles under both indicator models. Under Model A, the top two prevalence is higher in AV-only sessions ($\Delta p = 0.1313$, 95% CI [0.0245, 0.2332]). Whereas the contrast for the single highest risk profile is not distinguishable ($\Delta p = -0.0395$, 95% CI [−0.1239, 0.0411]). Under Model B, AV-only sessions show higher prevalence for both the highest risk profile ($\Delta p = 0.0809$, 95% CI [0.0182, 0.1412]) and the combined top two profiles ($\Delta p = 0.1456$, 95% CI [0.0788, 0.2100]). These findings suggest that AV-only conditions are associated with a larger share of sessions in high risk interaction regimes in Newcastle, implying smaller temporal safety margins for a non-trivial subset of encounters. This result aligns with previous research. Chen et al. Chen et al. (2020) found that pedestrians tend to accept larger gaps when interacting with AVs, potentially leading to more risky encounters if pedestrians misjudge the safety of the interaction. On the other hand, Velasco et al. Velasco et al. (2021) demonstrated that AVs sometimes induce more cautious behaviour in pedestrians, but this effect is context dependent. In some cases, pedestrians accepted smaller gaps when interacting with HDVs, highlighting that pedestrians' decision making in crossing situations varies depending on the vehicle type and the surrounding context. Our results indicate that, in Newcastle, AVs are perceived as riskier in terms of interaction criticality. This suggests that pedestrians may misinterpret AV intent or feel more uncertainty during AV encounters compared to HDVs.

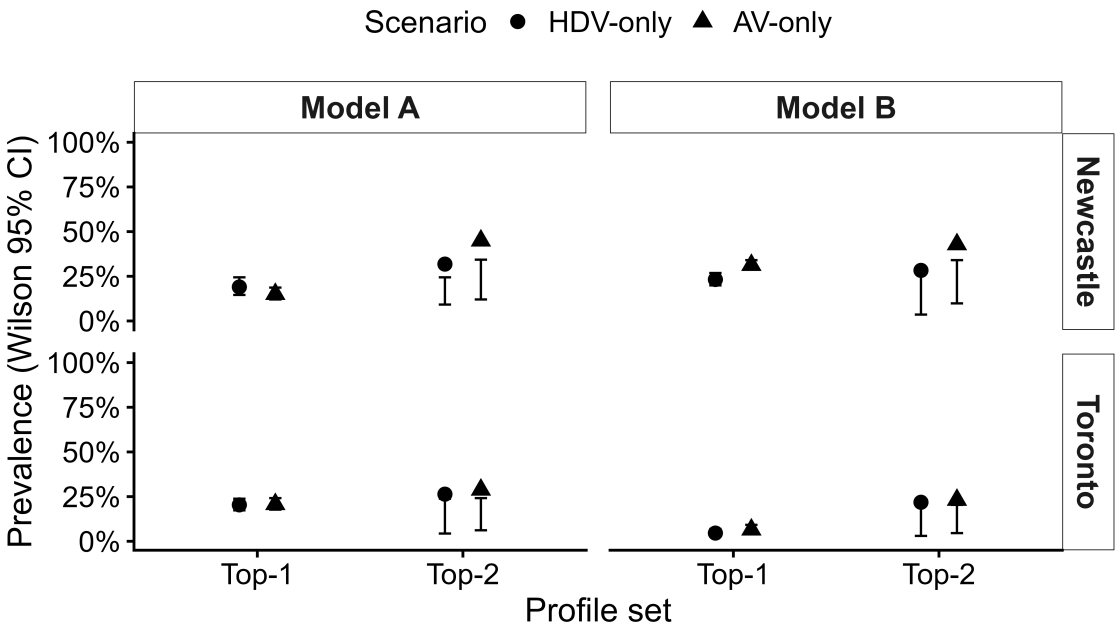


Figure 1: Prevalence of the highest risk end of the latent profile structure by scenario type.

In Toronto, AV-only versus HDV-only contrasts are small and not distinguishable for both indicator models. For Model A, $\Delta p = 0.0038$ for the Top-1 risk profile (95% CI $[-0.0611, 0.0685]$) and $\Delta p = 0.0246$ for the top two profiles (95% CI $[-0.0472, 0.0960]$). For Model B, $\Delta p = 0.0193$ for the highest risk profile (95% CI $[-0.0233, 0.0618]$) and $\Delta p = 0.0131$ for the top two profiles (95% CI $[-0.0638, 0.0899]$). Within the estimation samples used for each model, Toronto exhibits broadly similar prevalence of high-risk profiles across AV-only and HDV-only scenarios. However, non-distinguishable contrasts do not imply equivalence. They indicate that the available data do not rule out small differences in either direction. Differences between Model A and Model B are consistent with their indicator definitions. Model A relies on MTTC, which is defined only for collision-course events under the assumed motion model and therefore emphasizes imminent collision configurations. Model B relies on T_2 and CTTC, which remain informative in near-miss encounters because they characterize arrival timing and closing urgency even when a collision is not predicted. This distinction is consistent with the stronger AV-only versus HDV-only differences in Newcastle under Model B and the weaker separation observed for Toronto Model B in Table 1.

Risk profile prevalence and indicator structure

Figure 1 confirms the prevalence contrasts shown in Table 2. In Newcastle, AV-only sessions showed a higher prevalence of high-risk profiles, particularly under Model B. In Toronto, AV-only and HDV-only sessions showed similar prevalence across both models. Figure 2 shows that lower risk ranks were consistently associated with smaller temporal margins and smaller CTTC values. This supports the internal ordering of the latent profiles. Separation between adjacent ranks was clearer in Newcastle than in Toronto, especially for Model B. This overlap is consistent with the lower entropy observed for this specification and suggests that the T_2-based indicator space captures a more continuous spectrum of urgency in this context. To help in interpretation, we next describe what these ranks could represent as hypothesized pedestrian crossing stances expressed within a session.

Interpretation of risk ranks

Risk ranks summarize the urgency of the most critical moment within each session. Lower ranks correspond to lower temporal margins and stronger closing urgency. Higher ranks reflected more buffered encounters. In both locations and indicator sets, the latent profiles followed this general ordering although the degree of separation differed by model. In Newcastle, the rank structure was more distinct. This is particular in Model B where T_2 remained informative for near miss encounters. This suggests that the joint use of T_2 and CTTC captured a wider range of interaction regimes. Those regimes vary from high-urgency sessions with tight timing and strong closing to

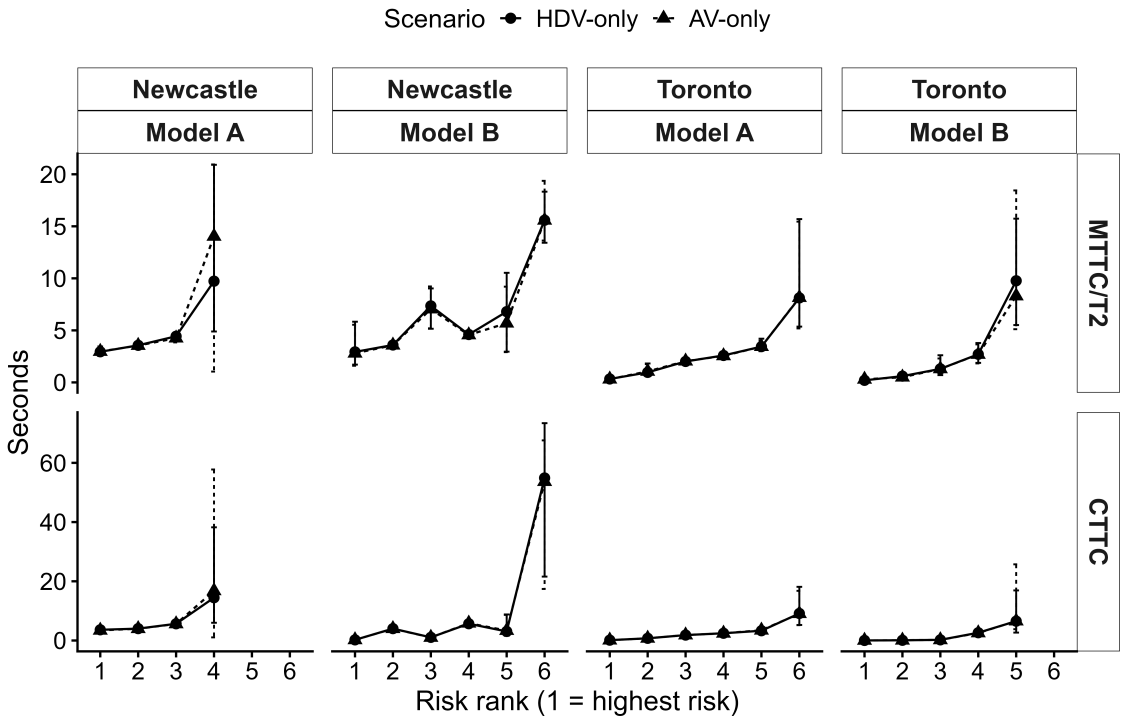


Figure 2: Indicator levels by risk rank showing medians with 10th–90th percentile ranges.

highly buffered sessions with large temporal margins. In Toronto, the same ordering was also present. However, adjacent ranks overlapped more specially in Model B. This is consistent with the weaker classification diagnostics. These ranks should be interpreted as descriptive summaries of interaction regimes rather than fixed behavioural categories. Their main value is that they reveal whether a larger share of sessions falls into high-urgency conditions. It provides a more informative safety comparison than relying only on average surrogate safety measures.

4 Conclusions

This paper applies LPA to SSMs to identify and compare latent pedestrian–vehicle interaction risk profiles across vehicle types and locations. This profile based approach ensures that rare but consequential, time critical sessions are not masked by aggregate summaries which preserves heterogeneity in interaction criticality. It also provides valuable insights for evaluating AV safety by identifying how AVs may affect pedestrian behaviour. Using immersive VR data from Toronto and Newcastle, each 60 s session is represented by its most critical time instance, defined by the session minimum of MTTC (Model A) or T_2 (Model B), paired with the corresponding CTTC at the same time step. Key results indicate that the Newcastle data shows clearer separation between latent risk profiles. In Toronto Model B, the separation is less clear likely due to the indicator structure. Additionally, in Newcastle, AV-only sessions led to more high risk interactions compared to HDV-only sessions. This suggests that AVs introduce higher levels of risk in some contexts. In contrast, Toronto did not show this trend, indicating that location and context play a significant role in the impact of AVs on pedestrian risk. The resulting profiles form an ordered risk rank that spans from high urgency (small temporal margins and urgent closing) to highly buffered encounters. These ranks likely reflect distinct pedestrian crossing stances, ranging from risk accepting to highly cautious. We can better interpret how interaction criticality varies across contexts through analyzing these profiles. Future work should explicitly model the influence of control variables embedded in the VR experimental design. These include traffic flow, weather, time of day, pedestrian density, social conformity, median presence, eHMI availability, and AV yielding behaviour. Including these factors into a joint modelling framework will allow assessment of how specific design and operational features shift sessions between latent risk profiles, rather than only comparing aggregate prevalence by vehicle type. Extending the framework to hybrid models that link latent profile membership to behavioural and contextual covariates would further support hypothesis driven evaluation of AV design choices and enable more direct translation of profile based findings into safety guidance for AV deployment.

Acknowledgements

This research is conducted under the project "Virtual Reality to study the role sOcIal Conformity in the acceptance of Autonomous vehicles (VeROnICA)," funded by the Economic and Social Research Council (United Kingdom).

References

Al-Haideri, R., Ismail, K., & Weiss, A. (2025). Beyond traditional models: a discrete choice approach to vehicle interaction modelling at roundabouts. *Transportmetrica A: Transport Science*, 1–45. doi: 10.1080/23249935.2025.2552982

Chen, W., Zhang, Y., Xu, W., Yang, Z., & Zhang, X. (2020). Comparison of pedestrians' gap acceptance behavior towards automated and human-driven vehicles. In *International conference on human-computer interaction* (pp. 132–144). Springer International Publishing. doi: 10.1007/978-3-030-50617-1_11

Hosseini Shoabjareh, A., Mamdoohi, A. R., & Nordfjærn, T. (2021). Analysis of pedestrians' behaviour: A segmentation approach based on latent variables. *Accident Analysis & Prevention*, *157*, 106160. doi: 10.1016/j.aap.2021.106160

Laureshyn, A., Svensson, Å., & Hydén, C. (2010). Evaluation of traffic safety, based on micro-level behavioural data: Theoretical framework and first implementation. *Accident Analysis & Prevention*, *42*(6), 1637–1646. doi: 10.1016/j.aap.2010.03.021

Laureshyn, A., & Várhelyi, A. (2018). *The swedish traffic conflict technique: Observer's manual*. Lund, Sweden.

Maji, A., & Ghosh, I. (2025). A copula-based multivariate extreme value framework for roundabout safety evaluation under mixed traffic. *Accident Analysis & Prevention*, *221*, 108219. doi: 10.1016/j.aap.2025.108219

Nagin, D. S., & Odgers, C. L. (2010). Group-based trajectory modeling in clinical research. *Annual Review of Clinical Psychology*, *6*, 109–138. doi: 10.1146/annurev.clinpsy.121208.131413

Nazemi, M., Rababah, B., Ramos, D., Zhao, T., & Farooq, B. (2025). Decoding pedestrian stress on urban streets using electrodermal activity monitoring in virtual immersive reality. *Transportation Research Part C: Emerging Technologies*, *171*, 104952.

Ozbay, K., Yang, H., Bartin, B., & Mudigonda, S. (2008). Derivation and validation of new simulation-based surrogate safety measure. *Transportation Research Record: Journal of the Transportation Research Board*(2083), 105–113. doi: 10.3141/2083-12

van Lissa, C. J., Rosenberg, J. M., Beymer, P. N., Anderson, D. J., & Schmidt, J. A. (2018). tidylpa: An R package to easily carry out latent profile analysis (lpa) using open-source or commercial software. *Journal of Open Source Software*, *3*(30), 978. doi: 10.21105/joss.00978

Velasco, J. P. N., Dufresne, T., Vinnitskaya, I., Cho, J., Murphey, J. T., Farooq, B., & Al-Haideri, R. (2021). Will pedestrians cross the road before an automated vehicle? the effect of drivers' attentiveness and presence on pedestrians' road crossing behavior. *Transportation Research Interdisciplinary Perspectives*, *12*, 100466. doi: 10.1016/j.trip.2021.100466