SUMO User Conference 2026
TIBconferencesession
TIB-OP will set DOI with \TIBdoi

Published: 1970-01-01

Investigating Traffic Effects of Control Transitions in Level 3 Conditional Driving with SUMO

Robert Alms[1], and Peter Wagner[12]

[1]Institute of Transportation Systems, German Aerospace Centre (DLR), Rutherfordstr. 2, 12489 Berlin, Germany

[2]Institute of Land and Sea Transport Systems, TU Berlin, Kaiserin-Augusta-Allee 104, 10553 Berlin, Germany

*Correspondence: Robert Alms

Abstract: Level 3 conditional automated driving relies on system-initiated takeovers when operating limits are reached. Even successful transitions of control (ToCs) can disturb traffic and accumulate in dense flow. Building on our prior studies, we present complementary SUMO motorway simulations and assess the impact of potential traffic-management measures that distribute takeover request timing. Unmanaged operation increases time headways, reduces capacity, and elevates surrogate safety events, with the strongest effects at intermediate mixed-autonomy shares. Managed timing mitigates these impacts, but ToC-related losses remain inherent to Level 3 operation.

Keywords: level 3 automation, transition of control (ToC), takeover request (ToR)

1. Introduction

Over the past decade, vehicle automation has advanced rapidly, yet deployed systems still depend on the human driver under clearly bounded operating conditions. Level 3 conditional driving, as specified on the SAE automation scale [1], is moving from pilot deployments toward market rollout [2]. Within its operational design domain (ODD), the vehicle may control the driving task, but it must request that the driver retake control within a minimum time span when the ODD boundary is approached or system limits are reached [3], [4]. These transitions of control (ToCs) expose a central automation challenge: drivers who have been disengaged must quickly rebuild situational awareness while simultaneously stabilizing manual vehicle control. Even non-emergency, system-initiated ToCs can transiently alter car-following behaviour and degrade control quality [5]. To ensure a safe takeover within a mandated lead time under UN R157 [6], the automated system may pre-emptively increase headway by reducing speed. In dense traffic, frequent ToCs can therefore introduce systematic disturbances that propagate upstream [7], degrading traffic efficiency and capacity and shaping safety-relevant interactions even when every takeover succeeds. SUMO [8] provides a framework to investigate such effects using microscopic simulation while assessing network-level traffic performance in mixed-autonomy conditions. The ToC device enables traffic simulations that combine modelling of individual ToC behaviour with traffic-level performance assessment to quantify such ToC-induced effects in dense mixed-autonomy traffic [9], [10]. Outside the EC project TransAID [11] and our own follow-up work, we are not aware of other simulation studies that explicitly investigate traffic-level effects of Level 3 ToCs in mixed-autonomy flow. Against this backdrop, this paper revisits key findings from our previous studies [12]–[14] and, within this scope, presents additional results (Section 4). Specifically, it exhibits time headway distributions and lane-change patterns to complement capacity-related findings in [14], and it adds safety-relevant metrics to further discuss network-level consequences of ToCs.

2. Control transitions in Level 3 conditional driving

Control transitions denote the bidirectional transfer of authority over the driving task between a human driver and the vehicle automation, constituted by a reallocation of longitudinal and/or lateral control authority and responsibility [15]. In this paper, transitions of control refer to automation-to-driver handovers in Level 3 conditional driving, where the system must prompt the driver to resume control when ODD boundaries are approached or system limits are reached [1]. Such prompts are issued as a takeover request (ToR) via the HMI, and their design can influence how quickly drivers return attention to the driving task [16], [17]. Since Level 3 permits engagement in non-driving related tasks, drivers may be less prepared at ToR, which can prolong takeover time [15], [18]–[20] and degrade post-takeover control [17], [21].

For these system-initiated, non-emergency transitions, the key timing relation is between the *available lead time* provided by the system[1] and the driver's *takeover time* required to re-establish situational awareness and stabilize manual vehicle control. If takeover time remains within the available lead time, the transition is considered successful, but may still involve a brief transient in performance [17]. Takeover time alone is not a sufficient indicator of quality, since rapid takeovers can still compromise post-transition stability [22], while takeover timing and control quality also depend on traffic and task demands. In particular, dense traffic has been shown to prolong takeover time and degrade post-takeover performance [5]. If takeover time exceeds the available lead time, the transition is unsuccessful and the system must initiate a fallback called a minimum risk manoeuvre (MRM) to reach a safe state. In addition, ToR lead time and trigger policy affect takeover timing [17], [21], while HMI modality and content influence post-transition performance [23], [24].

Existing ToC modelling is typically split between driver-centred takeover behaviour and automation-centred takeover requests and fallback manoeuvres, while an integrated representation of the full transition process is largely missing. SUMO addresses this gap with its ToC device [9], [10]. It enables switching between automated and manual driving by changing parameter sets in the underlying car-following models, or by substituting one model for another. The device represents the key phases of a transition, including an automated preparation phase after a takeover request and a temporary reduction in driver performance after control is resumed. Figure 1 summarises this logic as a state machine.

[1] Subject to regulatory requirements; currently $T_{\text{lead}} = 10\,\text{s}$ in UN R157 [6].

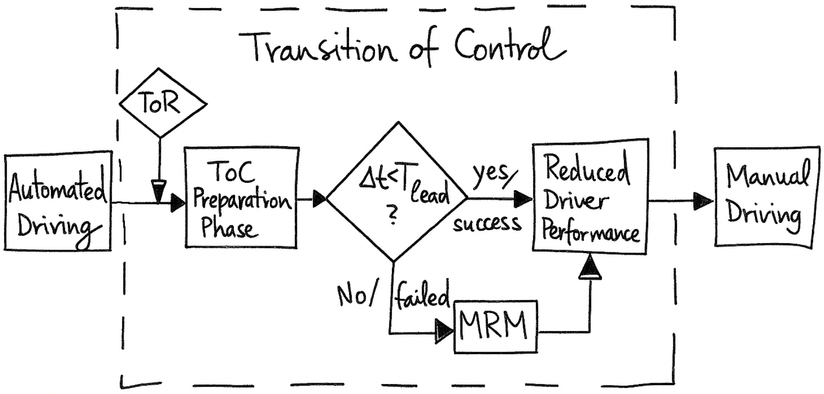

Figure 1. *State machine representation of the SUMO logic for transitions of control. In automated driving, a takeover request activates the preparation phase. A successful transition occurs if* $\Delta t < T_{\text{lead}}$ *and leads to reduced driver performance followed by manual driving. Otherwise, the transition is unsuccessful and the model enters the MRM state, which persists until driver intervention (or a minimum risk condition, typically standstill).*

In SUMO, a successful control transition is represented by two coupled mechanisms that shape the trajectory in Figure 2. First, the ToR triggers a ToC preparation phase in the automated vehicle (AV) during which the automation can reduce speed to build additional headway to the leading vehicle and thereby increase the available time budget for takeover [6], [9]. Second, once the driver resumes control, the model accounts for temporarily reduced driver performance and a recovery toward normal manual driving, which can manifest as a transient in speed and car-following behaviour after the handover [9].

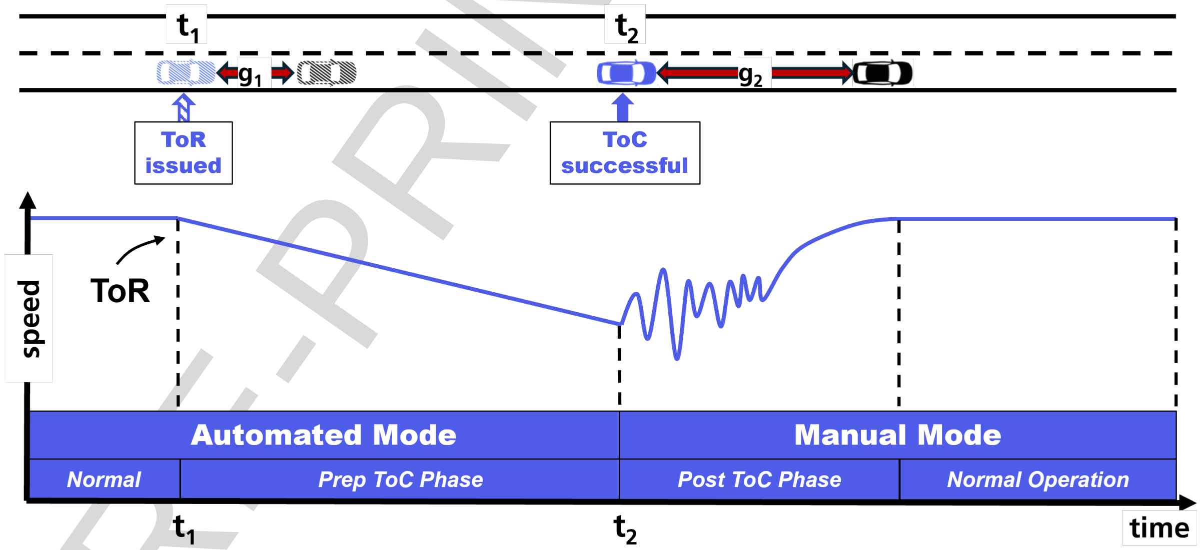

Figure 2. *Schematic timeline of a successful transition of control in a car-following situation. At* t_1*, a ToR initiates a preparation phase in automated mode, during which the automation reduces speed to increase the headway from* g_1 *to* g_2*. At* t_2*, manual control is resumed and a transient phase of reduced driver performance follows before returning to steady manual driving.*

Extending the single-vehicle ToC mechanism to a multi-vehicle setting, Figure 3 illustrates how the ToC preparation phase interacts with car-following dynamics in a platoon. The simplified single-lane experiments use identical vehicle parametrization and consider a 100% AV platoon in panel (a) and a 50/50 AV–MV mix in panel (b)[2]. Across both compositions, successive preparation phases introduce speed reductions that propagate upstream. When consecutive AVs enter the preparation phase while already embedded in disturbed flow, they start from lower speeds and altered spacing, which requires stronger and/or earlier deceleration to build the desired headway within the same lead-time constraint. As a result, the disturbance accumulates along the platoon, with differing severity between panels (a) and (b) while the characteristic upstream amplification persists in both cases.

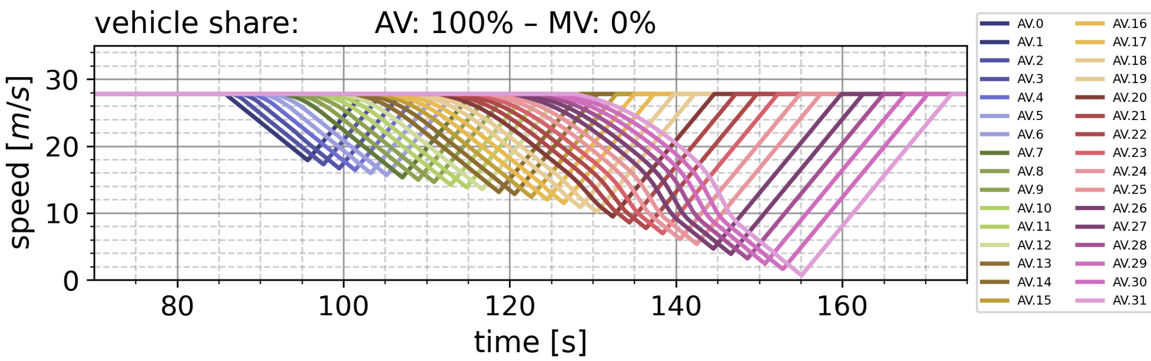

(a)

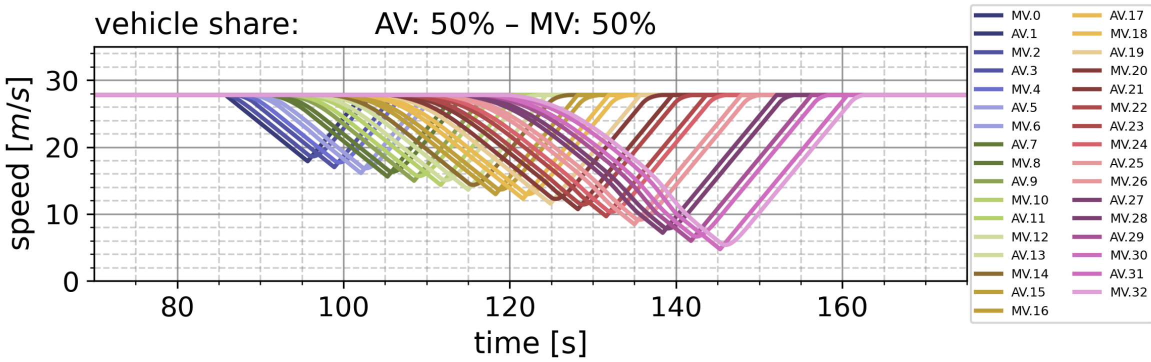

(b)

Figure 3. *Speed trajectories in simplified single-lane platoon experiments illustrating upstream accumulation of ToC-induced speed reductions during the preparation phase. Panel (a) shows a 100% AV platoon. Panel (b) shows a 50/50 AV–MV mix. In both cases, successive preparation phases introduce speed reductions that propagate upstream and become increasingly pronounced along the platoon, even when individual ToCs are successful.*

3. Simulation experiment setup

To translate the control transition mechanisms discussed above into a reproducible traffic-flow experiment, we implement a SUMO scenario on a real-world motorway network with a hypothetical localised ODD-limited zone that triggers system-initiated takeovers. Whereas the platoon experiments in Figure 3 illustrate ToC effects under idealised conditions, the baseline for the motorway flow was calibrated to real-world detector data. Level 3-capable AVs are then introduced into this calibrated stream to form mixed-autonomy scenarios. The ToC device is parameterised in line with R157.[3]

[2]Manual vehicles (MVs) use SUMO's default model; AVs are emulated using the ACC model [25], [26].
[3]For a detailed description of the replicated simulation setup and calibration, we refer to [14].

3.1. Scenario

Figure 4 provides an overview of the setup. Panel (a) shows a use case with a localised ODD boundary when vehicles approach a no-automated driving area, and the lower subpanel shows the end of the ODD zone in an exemplary SUMO snapshot with dedicated vehicle colours indicating different driving states. Panel (b) depicts the simulated network and the ODD-limited zone. Panel (c) summarises heterogeneity across vehicle classes by showing the distributions of four key behavioural variables.

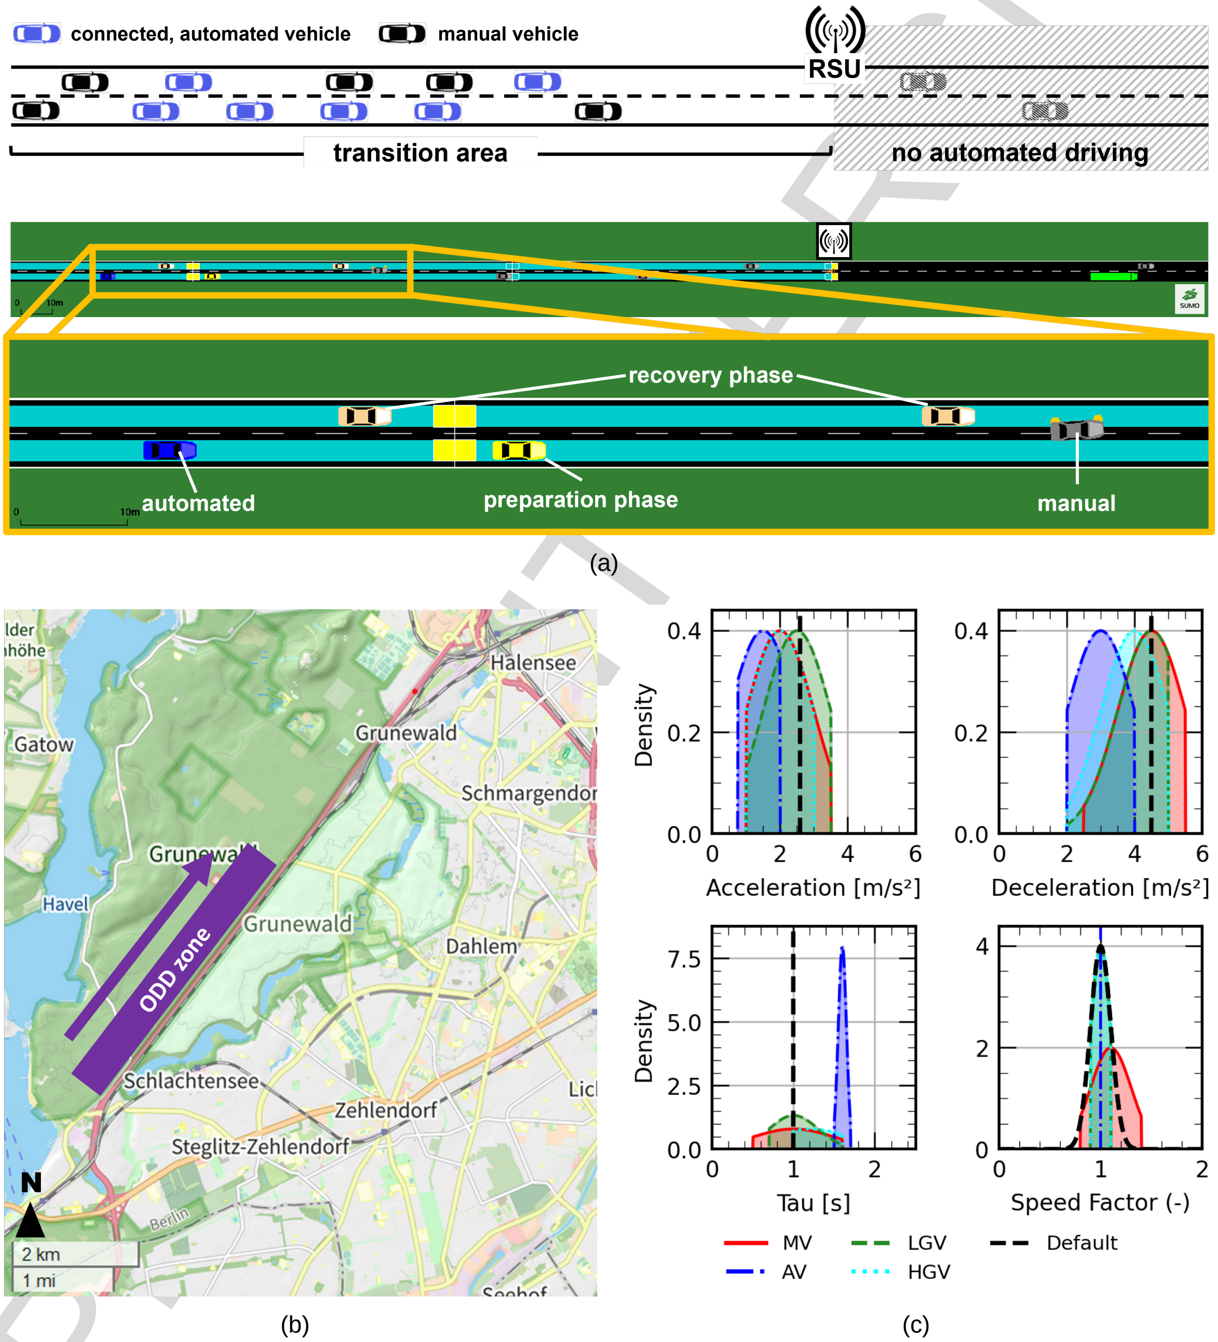

Figure 4. Simulation experiment setup. (a) ToR transmission along the transition area via an RSU, and SUMO GUI snapshot illustrating the end of the ODD-limited zone with vehicle colours indicating the driving state. (b) Simulated German motorway network with the ODD-limited zone. (c) Vehicle-class parameterisation shown as distributions for four key behavioural parameters. MV=Manual vehicle; AV=Automated vehicle; LGV/HGV=Light/Heavy good vehicle.

Figure 4, panel (a) illustrates a concept for the ToR transmission along the transition area via a roadside unit (RSU) to V2X-capable connected automated vehicles. This idea is introduced to enable a comparison between different approaches on handling ODD boundaries from a infrastructure and management perspective. A rather simple, but practical approach would be to mandate ToRs at the downstream ODD boundary, leading to ToCs clustering near the end of the ODD zone (unmanaged case). On the contrary, a more sophisticated, heuristic approach to dispatch the ToR timing is applied via the scheduling algorithm in [9] to facilitate a management of ToCs (managed case).

3.2. Replicating capacity reductions

As a first step, we replicated the capacity results reported in [14] to validate our implementation and to obtain consistent traffic demand levels for subsequent analyses. For each traffic composition with increasing AV share, we determine the maximum flow q in veh/h in the simulation for three cases: "No ToCs" serves as a baseline without control transitions. "Unmanaged" triggers ToCs at the end of the ODD zone. "Managed" distributes ToR timing for AVs within the transition area. Figure 5 summarises the replicated capacities across compositions.

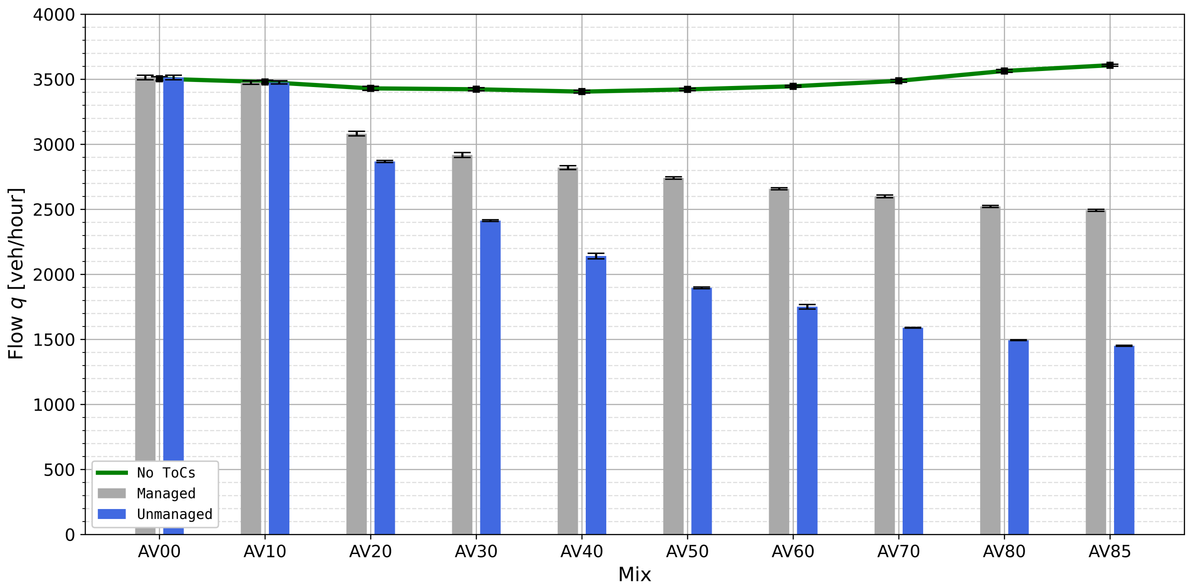

Figure 5. *Maximum flow q as a measure of capacity across traffic compositions with increasing AV share for the three ToC cases. Lines and bars indicate the mean across simulation seeds, and error bars indicate variability across seeds. The colour coding follows the legend.*

Compared to the no-ToC baseline, control transitions reduce capacity noticeably in the managed case, and even more so in the unmanaged case. Up to 10% AV share, capacities remain close to the baseline, whereas a clear drop becomes visible at 20% AV share. With further increases in AV share, the capacity loss grows, with the managed case mitigating some of the losses compared to the unmanaged case, which is most affected.

3.3. Further analyses and performance measures

For the subsequent analyses, we conduct ten simulations for each parameter combination (case and traffic mix) with different random seeds for stochastic vehicle insertion. For each composition and case, the input flow is set to the replicated capacity q from Figure 5. The vehicle type parametrisation scheme, including the ToC model parameters, follows [14, Table 2]. Driver takeover time after a ToR follows a skewed distribution with mean $7\,\mathrm{s}$, consistent with [12]. Under this configuration, fewer than 3% of ToCs across all runs exceed the regulated lead time $T_{\text{lead}} = 10\,\mathrm{s}$ and therefore trigger only brief MRMs of about 1–$3\,\mathrm{s}$.[4] Each simulation runs for 90 minutes, with a 30 min warm-up period to populate the network. All reported measures are computed over the remaining 60 min interval.

We assume AVs operate in Level 3 automated mode up to $100\,\mathrm{km/h}$ and remain automated until the end of the ODD zone. Vehicles are inserted upstream of the ODD zone in the northbound direction on one of two lanes and start in their respective initial driving mode, as shown in Figure 4(b).[5] They continue along the motorway segment and leave the network at the downstream boundary.

Within the ODD zone, we evaluate additional metrics. Time headway distributions are computed from vehicle trajectories, and lane-changing activity is derived from SUMO lane-change outputs. Safety-relevant interactions are quantified using surrogate safety measures (SSMs), specifically time-to-collision (TTC) [27] and modified deceleration rate to avoid a crash (MDRAC) [28], obtained via SUMO's dedicated SSM device.

4. Results

We next quantify the differences between managed and unmanaged operation across AV shares, with demand set to the replicated capacity levels.

4.1. Time headway distributions

Figure 6 shows time headway distributions across AV shares for the unmanaged (blue) and managed (gray) cases. With increasing AV penetration, headways shift to larger values and the distribution broadens, with the strongest changes occurring without management. When ToRs are scheduled, both the mean and the spread are reduced at a given AV share, consistent with the capacity trends noted in Section 3.2 (Figure 5).

4.2. Lane-changing effects

To characterise the ToC-related capacity losses, Figure 7 reports lane-change rates within the ODD zone as differences relative to the no-ToC baseline, expressed as events per $1000\,\mathrm{veh/h}$ at capacity. Panel (a) illustrates that excess lane-changing activity increases with AV share across traffic mixes. At low to mid AV shares from AV10 to AV40, the managed case exhibits higher rates than the unmanaged case, whereas from AV50 onward the ordering reverses and the difference increases with AV share. Panel (b) highlights that this activity is spatially structured. In managed runs, lane changes are concentrated near the ODD start and decay along the zone, while unmanaged runs show more dispersed activity with a build-up toward the ODD end. The

[4] MRMs with prolonged non-response can significantly disrupt traffic and are a central topic in current research on operationalisation and management, particularly in Level 4 automation. Yet, while MRMs are modelled in SUMO, they are not the focus of this study. Their impact is intentionally kept minor through the parametrisation to avoid obscuring the transitional impacts caused by the ToC preparation phase in more prevalent successful ToCs.

[5] HGVs are inserted on the right lane only.

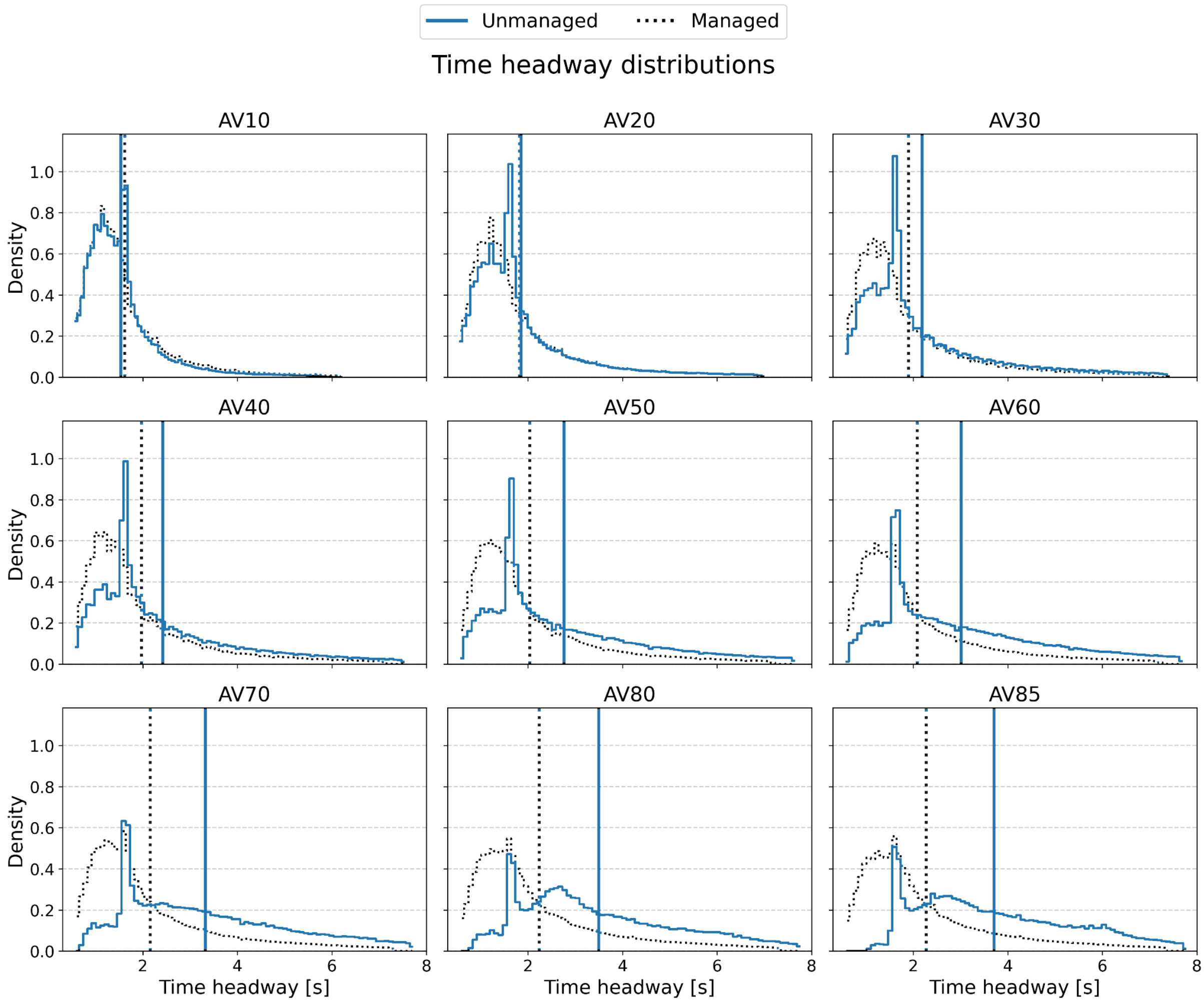

Figure 6. *Time headway distributions for traffic mixes AV10–AV85, comparing Unmanaged and Managed ToC operation. The mean of each distribution is indicated by a vertical line.*

spatial profiles are consistent with the aggregate rate differences in panel (a) for the same traffic mixes.

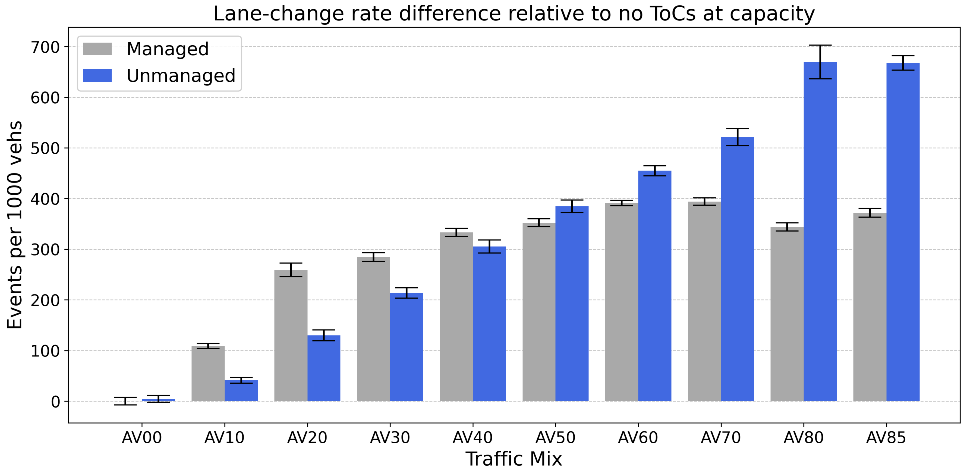

(a)

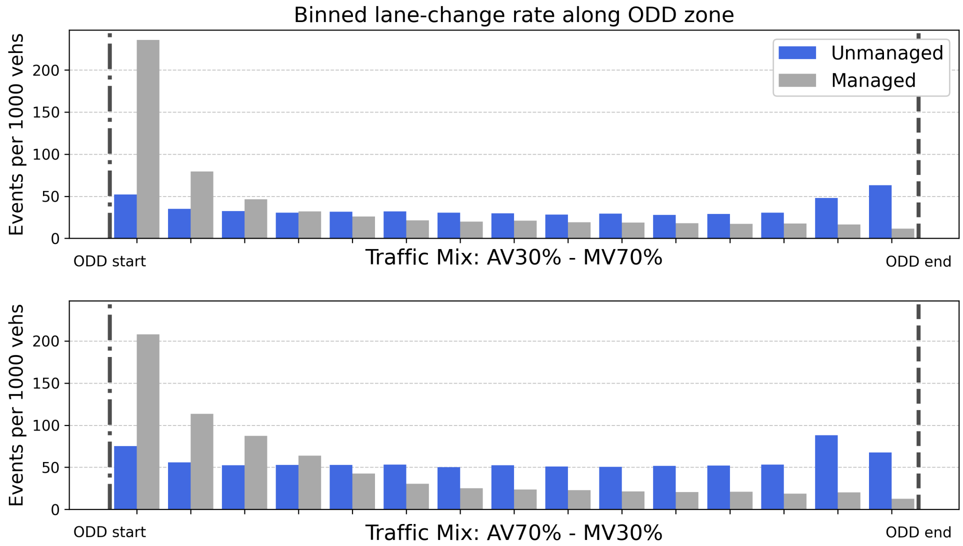

(b)

Figure 7. *Lane-change rates within the ODD zone for managed and unmanaged operation. Rates are reported as events per* 1000 *vehicles and normalised to the no-ToC baseline to correct for throughput differences. (a) Differences in lane-change rate at capacity as a function of traffic mix. Error bars denote the standard error across seeds. (b) Binned lane-change profiles along the ODD zone for AV30–MV70 and AV70–MV30. Vertical lines indicate the ODD boundaries.*

It is important to note that AVs do not change lanes during the ToC preparation phase. The observed patterns therefore primarily reflect reactive manoeuvres by surrounding vehicles in response to ToC-induced disturbances. The downstream concentration in the unmanaged case is consistent with more disruptive interactions near the capacity-critical ODD end and aligns with the higher mean headways and increased occurrence of large headways in Figure 6.

4.3. Safety impacts

Figure 8 summarises safety-relevant interactions in the ODD-limited motorway scenario using the SSM pair TTC and MDRAC. Event counts are reported per simulation hour and separated into inter-vehicle and intra-vehicle interactions. Overall, the unmanaged case exhibits consistently higher event rates than the managed case, with the clearest separation in the inter-vehicle panels. Across mixes, both TTC and MDRAC follow a non-monotonic pattern, reaching their highest levels at low-to-moderate AV penetration (around 20–40%) and decreasing toward higher AV shares. Inter-vehicle event rates exceed intra-vehicle event rates, suggesting that safety-relevant interactions are dominated by conflicts between different vehicle classes due ToC-induced disturbances.

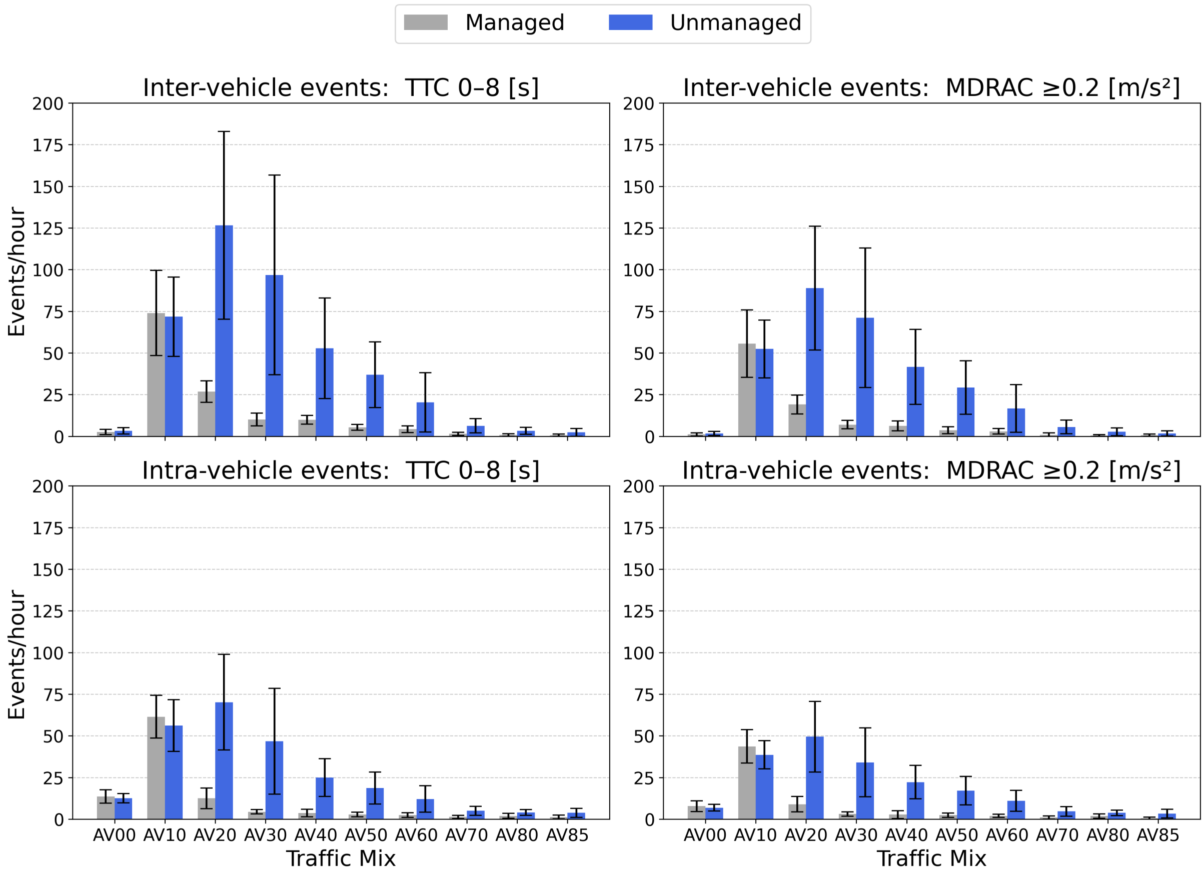

Figure 8. *SSM event frequency across traffic mixes as average events per hour for managed and unmanaged operation. Top panels show inter-vehicle conflicts; bottom panels show intra-vehicle conflicts. TTC events are counted for 0–8 s (left), and MDRAC events for* $\geq 0.2\,\mathrm{m/s^2}$ *(right). Both cases are evaluated at the same demand level, defined by the unmanaged case at* $v/c = 1$*, to avoid throughput-driven bias in event counts.*

5. Discussion

The additional results complement our previous findings on ToC-induced traffic impacts. We discuss three key points. First, our capacity analysis in [14] reported a mismatch at intermediate mixed-autonomy shares between capacity reductions predicted from simplified single-lane numerical experiments and the capacities measured in the full two-lane scenario (Figure 5). We attributed this underprediction in part to lane-changing interactions that are absent in the single-lane setting. The new results in Figures 6 and 7 provide evidence for this missing mechanism. They show that ToC-induced disturbances shift time headway distributions toward larger values, and that these disturbances are accompanied by systematic changes in lane-changing activity that depend on the ToR strategy. In the unmanaged case, the concentration of ToCs at the downstream ODD boundary is associated with higher lane-changing activity in its vicinity, inducing local disruptions that propagate upstream. In contrast, managed operation triggers ToRs earlier, shifting ToCs toward the beginning of the ODD zone and relocating the resulting reactive lane changes away from the downstream boundary, which mitigates disruptive impacts. Together, these patterns help explain why intermediate AV shares exhibit stronger capacity losses than predicted by single-lane experiments alone.

Second, our previous study [13] investigated ToC-related safety effects in a different ODD-limited use case with no Level 3 automated driving above $60\,\mathrm{km/h}$ and reported detrimental safety impacts with increasing AV shares. To provide a fuller picture of safety implications in mixed-autonomy traffic, we therefore evaluated SSMs for the use case considered in this paper (Section 3.1). The results indicate pronounced ToC-related safety impacts at low to mid AV shares (10–40%), which are substantially reduced under managed operation. Inter-vehicle event rates are consistently higher than intra-vehicle event rates by up to about 70%, suggesting that safety-relevant interactions are dominated by conflicts between different vehicle classes. This pattern is consistent with the lane-change findings, in that ToC-induced disturbances can trigger reactive lane-change manoeuvres and increased braking activity in surrounding traffic, particularly when ToCs concentrate near the downstream ODD boundary. At higher AV shares, safety benefits appear to counterbalance some ToC-induced impacts. The capacity recovery in the no-ToC baseline (cf. Figure 5 and the concurrent decline in SSM event rates are consistent with a smoother, more homogeneous flow governed by AV behaviour, in which ToC-induced disturbances and the associated reactive lane-change manoeuvres become less disruptive.

Third and last, our first study [12] investigated the traffic-efficiency performance of different ToR timing methods in a use case comparable to that of this paper. We subsequently extended this analysis by testing more recent algorithms, which performed notably better on the optimisation task. This also clarified that improving, or at least maintaining, traffic efficiency involves a trade-off between competing objectives once maximising automated driving distance is included as an additional objective. Overall traffic efficiency, consistent with the observed capacity reductions, is significantly reduced by ToCs. While traffic-management measures can alleviate some of these detrimental effects, they cannot restore baseline conditions. A short summary of the follow-up investigation is provided in Appendix A.

6. Conclusion

Transitions of control introduce disturbances into the traffic stream. While the local impact of an isolated transition may be minor, frequent transitions in dense traffic can accumulate, and their disruptive effects can propagate upstream even when every takeover succeeds. In our experiments, these effects arise from the preparatory gap-opening behaviour implemented to ensure safe takeover conditions, represented in SUMO's ToC device through an open-gap control mechanism. Although potential traffic-management measures can partially mitigate these adverse impacts, the underlying out-of-the-loop problem of automation inherent to Level 3 conditional driving with control transitions cannot be fully resolved.

Data availability statement

All data presented in this paper were generated using SUMO simulations (v1.25). The datasets will be made available upon request without undue reservation.

Author contributions

Conceptualization: RA & PW; data curation: RA; formal analysis: RA; funding acquisition: RA; investigation: RA; methodology: RA; project administration: RA; resources: RA; software: RA; supervision: PW; validation: RA; visualisation: RA; writing - original draft: RA; writing - review & editing: RA & PW.
All authors have read and agreed to the published version of the manuscript.

Competing interests

The authors declare that they have no competing interests.

References

[1] SAE International, "SAE International Recommended Practice: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles," SAE International, SAE Standard J3016_202104, Apr. 2021. DOI: 10.4271/J3016_202104.

[2] L. Hua, J. Antona-Makoshi, and L. Neurauter, "Status and challenges of level 3 automated driving systems," National Surface Transportation Safety Center for Excellence (NSTSCE), Virginia Tech Transportation Institute, Report 25-UT-172, Aug. 2025, Submitted: 2025-08-12. [Online]. Available: https://hdl.handle.net/10919/137487.

[3] British Standards Institution, "PAS 1883:2020 - Operational design domain (ODD) taxonomy for automated driving systems (ADS) – Specification," British Standards Institution, London, UK, Tech. Rep., 2020, Accessed: 2025-03-21. [Online]. Available: https://www.bsigroup.com/globalassets/localfiles/en-gb/cav/pas1883.pdf.

[4] "ASAM OpenODD Base Standard: Specification," ASAM e.V., Standard 1.0.0, 2025. [Online]. Available: https://publications.pages.asam.net/standards/ASAM_OpenODD/ASAM_OpenODD/latest/specification/index.html.

[5] C. Gold, M. Körber, D. Lechner, and K. Bengler, "Taking over control from highly automated vehicles in complex traffic situations: The role of traffic density," *Human Factors*, vol. 58, no. 4, pp. 642–652, 2016. DOI: 10.1177/0018720816634226.

[6] UNECE, "Addendum 156 – un regulation no. 157 - amendment 4 - uniform provisions concerning the approval of vehicles with regard to automated lane keeping systems," United Nations Economic commission for Europe, Tech. Rep., Mar. 2023, Accessed on Mar. 12th, 2025. [Online]. Available: https://unece.org/transport/documents/2023/03/standards/un-regulation-157-amend4.

[7] S. Maerivoet, L. Akkermans, K. Carlier, *et al.*, "Transaid deliverable 4.2 – preliminary simulation and assessment of enhanced traffic management measures," European Commission, Tech. Rep., Oct. 2020. DOI: 10.3030/723390. [Online]. Available: https://cordis.europa.eu/project/id/723390/results.

[8] P. Alvarez Lopez, M. Behrisch, L. Bieker-Walz, *et al.*, "Microscopic Traffic Simulation using SUMO," in *The 21st IEEE International Conference on Intelligent Transportation Systems*, IEEE, Nov. 2018, pp. 2575–2582. [Online]. Available: https://elib.dlr.de/127994/.

[9] L. Lücken, E. Mintsis, K. Porfyri, R. Alms, Y.-P. Flötteröd, and D. Koutras, "From Automated to Manual - Modeling Control Transitions with SUMO," in *SUMO User Conference 2019*, M. Weber, L. Bieker-Walz, R. Hilbrich, and M. Behrisch, Eds., ser. EPiC Series in Computing, vol. 62, EasyChair, 2019, pp. 124–144. DOI: 10.29007/sfgk.

[10] E. Mintsis, D. Koutras, K. Porfyri, *et al.*, "Transaid deliverable 3.1 - modelling, simulation and assessment of vehicle automations and automated vehicles' driver behaviour in mixed traffic," European Commission, Tech. Rep., Sep. 2020. DOI: 10.3030/723390. [Online]. Available: https://cordis.europa.eu/project/id/723390/results.

[11] TransAID, "Transition areas for infrastructure-assisted driving (transaid)," European Commission, H2020 Research Project 723390, 2021. DOI: 10.3030/723390. [Online]. Available: https://cordis.europa.eu/project/id/723390.

[12] R. Alms, A. Noulis, E. Mintsis, L. Lücken, and P. Wagner, "Reinforcement Learning-Based Traffic Control: Mitigating the Adverse Impacts of Control Transitions," *IEEE Open Journal of Intelligent Transportation Systems*, vol. 3, pp. 187–198, 2022. DOI: 10.1109/OJITS.2022.3158688. [Online]. Available: https://ieeexplore.ieee.org/document/9733028.

[13] R. Alms and P. Wagner, "Control Transitions in Level 3 Automation: Safety Implications in Mixed-Autonomy Traffic," *Safety*, vol. 10, no. 1, 2024, ISSN: 2313-576X. DOI: 10.3390/safety10010001. [Online]. Available: https://www.mdpi.com/2313-576X/10/1/1.

[14] R. Alms and P. Wagner, "Traffic capacity constraints from level 3 control transitions," *Frontiers in Future Transportation*, vol. Volume 6 - 2025, 2025, ISSN: 2673-5210. DOI: 10.3389/ffutr.2025.1600739. [Online]. Available: https://www.frontiersin.org/journals/future-transportation/articles/10.3389/ffutr.2025.1600739.

[15] Z. Lu, R. Happee, C. D. Cabrall, M. Kyriakidis, and J. C. de Winter, "Human factors of transitions in automated driving: A general framework and literature survey," *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 43, pp. 183–198, 2016, ISSN: 1369-8478. DOI: /10.1016/j.trf.2016.10.007.

[16] J. A. Miller, S. Nikan, and M. H. Zaki, "Navigating the handover: Reviewing takeover requests in level 3 autonomous vehicles," *IEEE Open Journal of Vehicular Technology*, vol. 5, pp. 1073–1087, 2024. DOI: 10.1109/OJVT.2024.3443630.

[17] B. W. Weaver and P. R. DeLucia, "A systematic review and meta-analysis of takeover performance during conditionally automated driving," *Human Factors*, vol. 64, no. 7, pp. 1227–1260, 2022. DOI: 10.1177/0018720820976476.

[18] L. Pipkorn, E. Tivesten, C. Flannagan, and M. Dozza, "Driver response to take-over requests in real traffic," *IEEE Transactions on Human-Machine Systems*, vol. 53, no. 5, pp. 823–833, 2023. DOI: 10.1109/THMS.2023.3304003.

[19] L. Pipkorn, M. Dozza, and E. Tivesten, "Driver visual attention before and after takeover requests during automated driving on public roads," *Human Factors*, vol. 66, no. 2, pp. 336–347, 2024. DOI: 10.1177/00187208221093863.

[20] E. Shi and K. Bengler, "Non-driving related tasks' effects on takeover and manual driving behavior in a real driving setting: A differentiation approach based on task switching and modality shifting," *Accident Analysis & Prevention*, vol. 178, p. 106 844, 2022, ISSN: 0001-4575. DOI: /10.1016/j.aap.2022.106844.

[21] A. D. McDonald, H. Alambeigi, J. Engström, *et al.*, "Toward computational simulations of behavior during automated driving takeovers: A review of the empirical and modeling literatures," *Human factors*, vol. 61, no. 4, pp. 642–688, 2019. DOI: 10.1177/0018720819829572.

[22] T. Gruden, S. Tomažič, and G. Jakus, "Post-takeover proficiency in conditionally automated driving: Understanding stabilization time with driving and physiological signals," *Sensors*, vol. 24, no. 10, 2024, ISSN: 1424-8220. DOI: 10.3390/s24103193.

[23] W. Morales-Alvarez, O. Sipele, R. Léberon, H. H. Tadjine, and C. Olaverri-Monreal, "Automated driving: A literature review of the take over request in conditional automation," *Electronics*, vol. 9, no. 12, 2020, ISSN: 2079-9292. DOI: 10.3390/electronics9122087.

[24] Q. Li, Z. Wang, W. Wang, *et al.*, "Systematically modeling take-over performance: Considering the indirect effect of meteorological visibility mediated by drivers' attention," *Accident Analysis & Prevention*, vol. 220, p. 108 174, 2025, ISSN: 0001-4575. DOI: /10.1016/j.aap.2025.108174.

[25] L. Xiao, M. Wang, and B. van Arem, "Realistic car-following models for microscopic simulation of adaptive and cooperative adaptive cruise control vehicles," *Transportation Research Record*, vol. 2623, no. 1, pp. 1–9, 2017. DOI: 10.3141/2623-01. [Online]. Available: https://journals.sagepub.com/doi/10.3141/2623-01.

[26] K. N. Porfyri, E. Mintsis, and E. Mitsakis, "Assessment of ACC and CACC systems using SUMO," in *SUMO 2018- Simulating Autonomous and Intermodal Transport Systems*, E. Wießner, L. Lücken, R. Hilbrich, *et al.*, Eds., ser. EPiC Series in Engineering, vol. 2, EasyChair, 2018, pp. 82–93. DOI: 10.29007/r343. [Online]. Available: /publications/paper/HvLp.

[27] J. C. Hayward, "Near-miss determination through use of a scale of danger," *Highway Research Record*, 1972. [Online]. Available: https://api.semanticscholar.org/CorpusID:53468218.

[28] Y. Kuang, X. Qu, J. Weng, and A. Etemad-Shahidi, "How does the driver's perception reaction time affect the performances of crash surrogate measures?" *PLOS ONE*, vol. 10, no. 9, pp. 1–13, Sep. 2015. DOI: 10.1371/journal.pone.0138617.

[29] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, "Stable-baselines3: Reliable reinforcement learning implementations," *Journal of Machine Learning Research*, vol. 22, no. 268, pp. 1–8, 2021. [Online]. Available: http://jmlr.org/papers/v22/20-1364.html.

[30] S. Fujimoto, H. van Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," *CoRR*, vol. abs/1802.09477, 2018. arXiv: 1802.09477. [Online]. Available: http://arxiv.org/abs/1802.09477.

[31] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *CoRR*, vol. abs/1707.06347, 2017. arXiv: 1707.06347. [Online]. Available: http://arxiv.org/abs/1707.06347.

A. Appendix

Figure 9 shows additional data points from follow-up reinforcement learning (RL) training runs based on the experiments originally presented in [12]. This study investigated the performance of different traffic management methods for ToR timing in an ODD zone, comparing heuristic approaches with RL models. Each data point in the figure corresponds to a trained policy or a heuristic solution, evaluated in a use case that follows the setup of this paper but uses a slightly different traffic composition and corridor length. Training was performed with an updated version of Stable Baselines3 [29], deploying TD3 [30] and PPO [31] algorithms. Across managed operating points, average travel time differs by up to about 25%, indicating substantial efficiency differences associated with the distance travelled in automated mode. The RL task turned out to be a multi-objective optimisation problem with competing objectives, namely maintaining automated driving operation within the ODD zone and ensuring sufficient traffic efficiency.

Overall, the results highlight a trade-off between traffic efficiency and automated driving distance when ToR timing is managed within an ODD-limited zone.

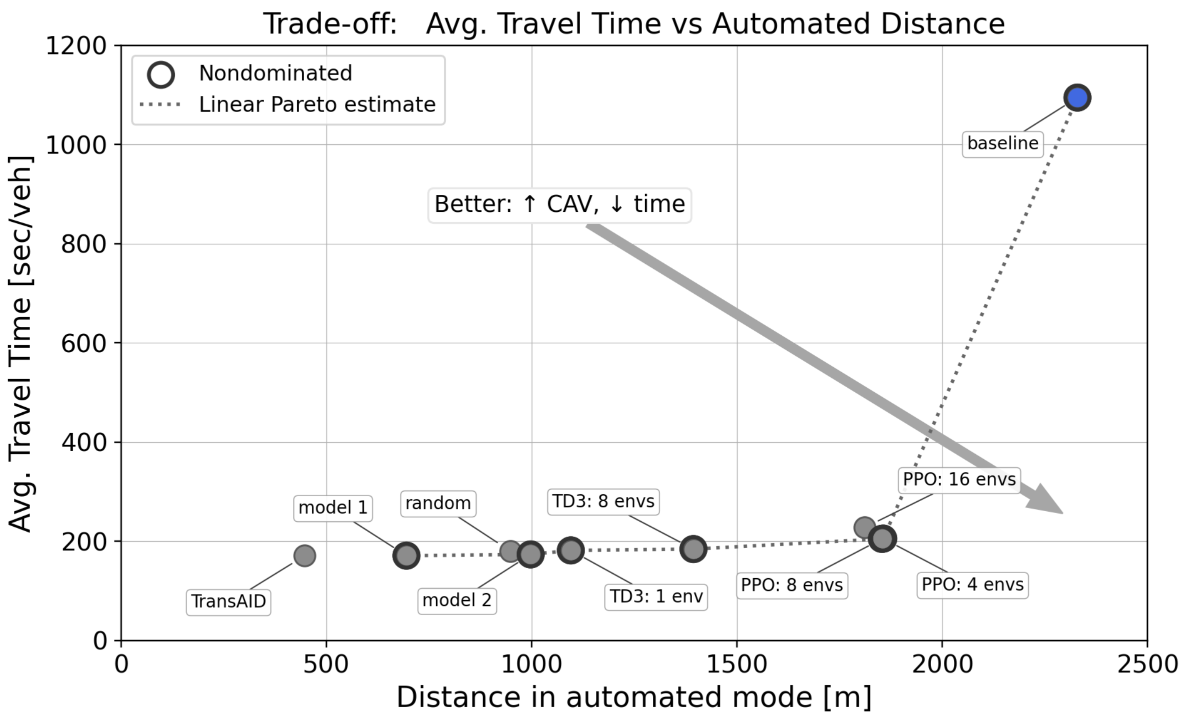

Figure 9. *Average travel time versus automated driving distance, illustrating the performance trade-off in managing ToR timing. Outlined markers denote non-dominated samples, and the dotted line provides a linear approximation of the Pareto front. The arrow indicates the preferred direction of improvement, with lower travel time and longer automated distance. The baseline point for the unmanaged case lies in a post-breakdown traffic regime and is therefore not directly comparable to the other data points for managed operation; it is shown for context, and the dotted line is extrapolated toward this point.*