Deutsches Zentrum
für Luft- und Raumfahrt
German Aerospace Center

Technische Universität Berlin

From Automated to Manual: Traffic Effects of Control Transitions in Level 3 Conditional Driving

vorgelegt von

Dipl.-Ing. Robert Alms
ORCID: 0000-0001-9950-3596

an der Fakultät V - Verkehrs- und Maschinensysteme
der Technischen Universität Berlin
zur Erlangung des akademischen Grades

Doktor der Ingenieurwissenschaften
– Dr.-Ing. –

genehmigte Dissertation

Promotionsausschuss:

Vorsitz:	Prof. Steffen Müller
Gutachter:	Prof. Kai Nagel
	Prof. Peter Wagner
	Prof. Peter Vortisch

Tag der wissenschaftlichen Aussprache: 4. Mai 2026

Berlin 2026

Für meine Eltern

Abstract

Conditional automated driving at Level 3 marks a decisive step on the path from manual to fully automated driving. Under defined operating conditions, the driver may divert attention from the traffic situation and hand over the driving task entirely to vehicle automation. When the system reaches the limits of its operating range, it prompts the driver to retake control within a specified time window. This takeover, termed a transition of control (ToC), exposes a central automation problem. A driver who has been out of the control loop for an extended period must, with reduced situational awareness, rapidly re-establish full vehicle control, make decisions and, where necessary, respond under time pressure. Such takeovers are particularly demanding in dynamic traffic, since even well-prepared, non-emergency ToCs can temporarily impair vehicle behaviour and the quality of manual control. To account for this, the automation takes preparatory measures by moderately increasing the time gap to the preceding vehicle during ToCs in dense traffic. Even without any failed takeovers, a high frequency of such ToCs can disturb traffic flow. This thesis therefore examines the resulting effects and their implications using microscopic traffic simulations and shows how system-initiated ToCs at the limits of Level 3 automation aggregate from single events into noticeable traffic-level effects. Several studies consider efficiency, safety and capacity. The focus is on the mechanisms associated with the sequence, timing and location of takeover triggers, as well as automation behaviour during ToCs. The results show that ToCs occurring in clusters cause a disproportionate disturbance of traffic flow up to complete breakdown; they degrade safety metrics and reduce capacity. The thesis also investigates traffic-management and V2X-enabled measures that can partially mitigate the consequences by distributing ToCs in time and space, but cannot fully restore baseline conditions of undisturbed, takeover-free traffic.

Kurzfassung

Das bedingt automatisierte Fahren der Stufe 3 markiert einen entscheidenden Zwischenschritt auf dem Weg vom manuellen zum vollautomatisierten Fahren. Unter definierten Betriebsbedingungen darf der Fahrer seine Aufmerksamkeit vom Verkehrsgeschehen abwenden und die Fahraufgabe vollständig an die Fahrzeugautomation übergeben. Erreicht das System jedoch die Grenzen seines Betriebsbereichs, fordert es den Fahrer innerhalb eines festgelegten Zeitfensters zur Rückübernahme der Kontrolle auf. Diese Übergabe der Kontrolle, genannt Transition of Control (ToC), legt ein zentrales Automatisierungsproblem offen: Ein Fahrer, der längere Zeit aus der Regelschleife war, muss mit reduziertem Situationsbewusstsein in kurzer Zeit wieder vollständige Fahrzeugkontrolle erlangen, Entscheidungen treffen und gegebenenfalls zeitkritisch reagieren. Solche Kontrollübergänge sind unter dynamischen Verkehrsbedingungen besonders anspruchsvoll, da selbst gut vorbereitete, nicht-notfallhafte ToCs das Fahrverhalten und die Qualität der manuellen Kontrolle kurzfristig beeinträchtigen können. Um diesem Phänomen Rechnung zu tragen, trifft die Fahrzeugautomation vorbereitende Maßnahmen, indem sie in dichtem Verkehr die Zeitlücke zum Vorderfahrzeug während des Übergabevorgangs moderat vergrößert. Auch ohne fehlgeschlagene Übergaben kann eine Häufung solcher ToCs den Verkehrsfluss auf diese Art stören. Diese Arbeit untersucht daher die daraus entstehenden Effekte und ihre Auswirkungen mithilfe mikroskopischer Verkehrssimulationen und zeigt, wie systeminitiierte ToCs an den Grenzen des Betriebsbereichs der Stufe-3-Automatisierung sich von Einzelereignissen zu spürbaren Netzeffekten aufsummieren. In mehreren Studien werden Effizienz, Sicherheit und Kapazität betrachtet. Der Fokus liegt dabei insbesondere auf der Analyse von Wirkmechanismen bezüglich der Aufeinanderfolge, des Zeitpunkts und des Ortes der Übergabe-Trigger sowie des Automationsverhaltens während der Übergaben. Im Ergebnis zeigen vor allem in Clustern auftretende ToCs eine überproportionale Störung des Verkehrsflusses bis hin zum vollständigen Zusammenbruch. Sie führen zu verschlechterten Sicherheitskenngrößen und senken die Kapazität. Untersucht werden ebenfalls V2X-gestützte Verkehrsmanagementmaßnahmen, die die Folgen durch zeitliche und räumliche Verteilung von ToCs zum Teil abmildern können, den Ausgangszustand eines ungestörten, übergabefreien Verkehrs jedoch nicht vollständig wiederherstellen.

Contents

List of Figures

List of Abbreviations

ACC	Adaptive Cruise Control
ALKS	Automated Lane Keeping System
AV	Automated Vehicle
CAV	Connected and Automated Vehicle
DDT	Dynamic Driving Task
DRAC	Deceleration Rate to Avoid a Crash
HGV	Heavy Goods Vehicle
LC	Lane Change
LGV	Light Goods Vehicle
MDP	Markov Decision Process
MDRAC	Modified Deceleration Rate to Avoid a Crash
MRM	Minimum Risk Manoeuvre
MV	Manual Vehicle
NDRT	Non-Driving-Related Task
No-AD	No Automated Driving
ODD	Operational Design Domain
PPO	Proximal Policy Optimisation
PRT	Perception-Reaction Time
RL	Reinforcement Learning
RSU	Roadside Unit
SB2	Stable Baselines 2
SB3	Stable Baselines 3
SSM	Surrogate Safety Measure

TB	Time Budget
TD3	Twin Delayed Deep Deterministic Policy Gradient
TM	Traffic Management
ToC	Transition of Control
ToR	Takeover Request
TTC	Time to Collision
V2X	Vehicle-to-Everything

1 Introduction

Automated driving has been part of the public imagination for decades and has long carried expectations of comfortable, safe, and effortless travel within an efficient transport system. The 1950s illustration in Figure 1.1 captures that optimism, showing a family relaxing while the car drives itself automatically: no traffic jams, no collisions, no driver fatigue. Each of these elements maps to a distinct research domain, such as engineering, traffic science, and human sciences, bringing a plethora of questions around the still-unfulfilled promise of fully autonomous driving.

Figure 1.1: Detail from "Driverless Car of the Future," advertisement for America's Electric Light and Power Companies, *The Saturday Evening Post*, 1956. Illustration provided by Curtis Licensing.

Yet after decades without a decisive breakthrough, vehicle automation has gathered renewed momentum over the past ten years, fuelled by technological progress and increased investment. Today the reality is more nuanced than those early visions: automation taxonomies specify how and when human drivers are involved in the driving task or, under defined conditions, may cede responsibility to an automated system. So-called "conditional" automation marks a critical shift: the system may manage the driving task within defined conditions, but must prompt the human to resume control within an adequate time frame when necessary. These transitions of control (ToC), particularly from automated to manual driving, exemplify the familiar "out-of-the-loop" conundrum in automation: drivers who have been out of the loop must make rapid decisions with reduced situational awareness while simultaneously re-establishing control. Such control transitions are, however, especially demanding under dynamic, time-critical conditions. Even well-prepared, non-emergency ToCs can alter car-following behaviour and briefly degrade manual control quality.

In addition to these human factors, "conditional" automation introduces a further complication, though: in denser traffic the automated system must ensure a safe ToC within a mandated minimum time span. It does so by pre-emptively increasing the time gap to the preceding vehicle by reducing speed. Therefore, when such control transitions occur more frequently, a succession of ToCs can lead to progressively larger fluctuations in upstream traffic flow. The resulting network-level effects can exert a measurable influence on overall traffic performance, even in the absence of any single failed takeover event. These ToC-induced effects are easy to underestimate, however, because most research focuses on the causes and immediate consequences of, or the subsequent management of, individual failed takeover events, from both a human-centred and an automation perspective. A traffic-level perspective that merges both ends of the spectrum into a cohesive analytical lens that provides a framework for assessing ToC implications in increasingly automated environments is largely absent. Given that the timing and location of such transitions could technically also be influenced via traffic management measures, it is natural to ask whether such measures can alleviate their traffic-level impacts in mixed-autonomy settings. Against this backdrop, the following research questions are posed:

RQ 1: How, and to what extent, do transitions of control affect traffic efficiency in dense traffic, and can traffic management measures alleviate these impacts?

RQ 2: Under what traffic and operational conditions do transitions of control lead to safety-relevant effects, and can mitigation measures reduce these effects?

RQ 3: What capacity constraints do transitions of control introduce, and are they recoverable through traffic management measures?

Direct, large-scale evidence on conditional driving remains scarce, as deployments are limited and suitable traffic-level datasets are not yet available. Accordingly, this thesis investigates network-level impacts using microscopic simulation, with ToCs represented explicitly by a dedicated model and embedded in scenarios designed to reveal their effects across a range of automation-penetration levels. In the absence of established traffic-management procedures, candidate measures are designed and compared with unmanaged baselines. The findings are presented in three publications (Papers I, II, and III), complemented by additional results that provide a comprehensive basis for addressing RQ 1 to RQ 3.

The remainder of the thesis is organised as follows: Chapter 2 conceptualises control transitions in Level 3 automation, summarises the empirical background, reviews and classifies existing models, specifies how ToCs are represented in microscopic traffic simulation, and defines the scope of the traffic-level effects considered. Chapter 3 assembles the associated publications and sets out key findings that provide the basis for the three main papers. Chapter 4 synthesises the studies, offers methodological reflections on the management of ToCs, considers safety at operational boundaries, examines capacity under ToC-induced constraints, and summarises findings and limitations. Chapter 5 offers a concise conclusion.

2 Control Transitions in Level 3 Automation

2.1 Conceptualising Control Transitions

Transitions of control in automated driving systems refer to the bidirectional shift of authority over the driving task between the human driver and the vehicle's automation, constituted in a reallocation of the longitudinal and/or lateral control task, as stated by [40]. The authors specify that control transitions delineate the full process of switching between so-called static driving states, which describe manual, automated, or shared driving tasks in such a driver-automation system. The driving task itself encompasses lateral and longitudinal vehicle control as well as continuous environmental monitoring, with transitions representing shifts in responsibility across one or more of these interdependent functions. These subtasks collectively form what the Society of Automotive Engineers (SAE) defines as the *Dynamic Driving Task (DDT)*, which underpins their classification of automation levels from Level 0 (no driving automation) to Level 5 (full driving automation) [68]. This classification incorporates the concept of an *Operational Design Domain (ODD)* [8], which specifies the environmental, geographic, and operational conditions under which an automated system is intended to function.

A classification of control transitions, also introduced by [40], differentiates by initiator (driver vs. automation) and control recipient (driver vs. automation), yielding four types: automation-initiated driver-in-control (AIDC), automation-initiated automation-in-control, driver-initiated driver-in-control, and driver-initiated automation-in-control. Within this framework, AIDC transitions, where the automation requests the human to resume full control of the DDT, have been among the most extensively studied and are often considered safety-relevant, especially when the driver is out of the control loop or engaged in non-driving related tasks (NDRT). Specifically in Level 3, the system performs the entire DDT within its ODD, and must initiate a transition of control to the human driver, typically via a *takeover request* (ToR), if the conditions of the current ODD are no longer met. This distinguishes it from

Level 2, where continuous driver supervision is required, and from Level 4, where the system manages fallback behaviour independently, including bringing the vehicle to a minimal risk condition without driver intervention. ToCs vary in urgency depending on the nature of the triggering event. *Non-emergency transitions* occur predictably, for example, when the system anticipates reaching the limits of its ODD, and typically allow the driver adequate time to resume control after a ToR. In contrast, *emergency transitions* result from sudden, unforeseen conditions such as system failures, sensor degradation, or unexpected obstacles, often requiring immediate human response. While both Levels 2 and 3 may involve emergency transitions, only Level 3 systems are responsible for detecting ODD boundaries and issuing ToRs under non-emergency conditions. In Level 2, the driver must continuously monitor the environment and respond without system prompting, effectively collapsing the distinction between emergency and non-emergency transitions from the driver's perspective.

When non-emergency transitions are initiated by a ToR in Level 3 automation, their success depends critically on the relationship between the system-provided *available lead time* and the driver's *takeover time*, that is, the time required to regain situational awareness and resume control. If the driver's takeover time remains within the available lead time, the result is a *successful transition*, potentially involving a brief phase of reduced performance before returning to stable manual operation. However, if the response exceeds this time budget or fails altogether, the transition is considered unsuccessful. In such cases, the system must initiate a fallback manoeuvre, called a *Minimum Risk Manoeuvre* (MRM), to bring the vehicle to a safe stop. Figure 2.1 illustrates this process by contrasting successful and failed transitions initiated by takeover requests.

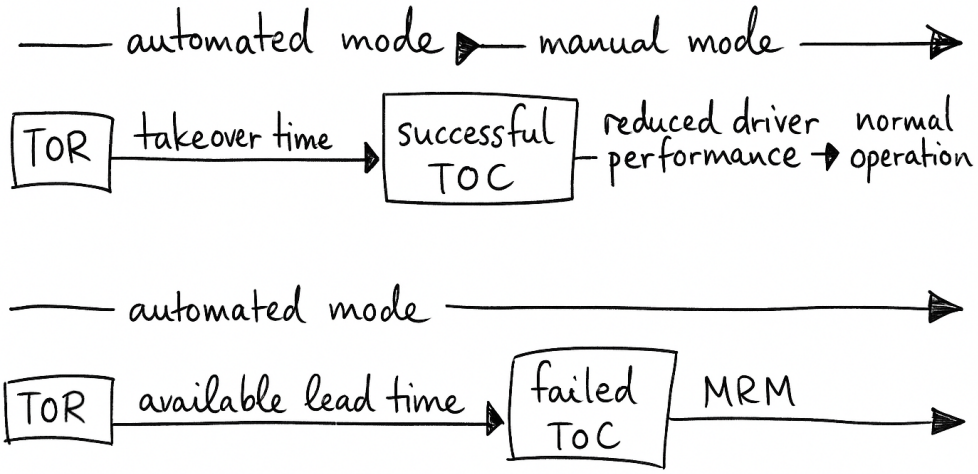

Figure 2.1: Conceptual model of a successful and failed transition of control (ToC) following a takeover request (ToR) in Level 3 automation.

2.2 Empirical Background of Control Transitions

Comprehensive and informative reviews by [40] and [45] synthesise findings from simulator, test-track, and on-road studies on Level 2 and 3 control transitions. They cover both driver- and system-initiated cases and report variations in takeover performance across lead times, ToR modalities, driver states, and environmental conditions. More recently, [30] extended the scope to Levels 0–4 with an engineering-centred focus, analysing key challenges, classification schemes, takeover experiments, performance evaluation methods, and authority allocation strategies, with particular emphasis on fallback manoeuvre planning and MRMs. With ToRs being the primary trigger for automation-initiated transitions in Level 3, a systematic review and meta-analysis by [86], along with concise reviews by [48] and [52], quantify the effects of lead time, secondary task engagement, and scenario complexity on takeover duration and quality, and examine timing, modality design, and the interaction of driver readiness with scenario demands. These reviews, often with a human-factors focus, highlight benefits of multimodal cues, particularly auditory and visual combinations, clear, context-rich ToR content, and adaptive delivery informed by driver monitoring.

Although most of these surveys are human-factors-centred, together they identify a wide range of aspects, ranging from driver states and situational awareness to automation capabilities such as sensing performance and fallback manoeuvre execution, the latter being covered to a more limited extent. A plethora of driving-simulator studies over the past decade on automation and out-of-the-loop phenomena have continued to inform questions of safe and efficient vehicle-automation design for ToCs. These aspects can be grouped into two categories: *human-centred factors* and *system-related factors*. A selection of relevant factors is examined in more detail below, within this categorisation.

2.2.1 Human-centred factors

Driver readiness and engagement

Driver readiness before a takeover request is a critical determinant of takeover performance, particularly in Level 3 conditional driving, since the human driver is allowed to engage in NDRTs. As reviewed by [40, 45], empirical studies have shown that situational awareness, trust, and anticipation substantially influence takeover time and quality. Engagement in NDRTs is consistently associated with degraded takeover performance [86, 72]. In a motion-base driving simulator study e.g. by [75], the authors found that visually demanding NDRTs delayed takeover initiation. Another simulator study by [16] reported that NDRTs delayed manual driving initiation, prolonged stabilisation, and increased lane-position variability,

with stronger effects for visually and cognitively demanding tasks. Although their experiments emulated Level 2 automation, they infer that human drivers assume Level 3 automation capability even for lower automation levels and therefore lack situational awareness, which indicates automation complacency. Recent public-road studies [63, 62] with Level 3 vehicles showed that NDRTs increased off-road glance durations, delayed gaze return to the forward roadway, and reduced mirror checking; high-visual-demand tasks exacerbated these effects. Response times increased notably for DDTs with NDRTs (median $\approx$ 5 s) compared with on-path gazing at ToR triggering [63]. Prior training can partly mitigate these effects: in a public-road study on German motorways, [55] found that targeted training shortened stabilisation time and improved control smoothness after voluntary takeovers. Beyond these examples, human-centred factors such as age and driving experience, trust calibration and prior exposure, vigilance/fatigue and workload tolerance, and the modality/intensity of NDRTs (visual, cognitive, manual) also influence takeover latency and post-takeover stabilisation [40, 45, 86].

Takeover time and performance quality

Even in non-critical situations, takeover time is a key indicator of driver readiness [70]. Again, NDRTs play a significant role in modulating such ToC properties. For example, [21] reported results from a simulator study for non-critical ToCs with median takeover times of 4.56 s without NDRTs, which increase significantly when secondary tasks were present (median 6.06 s; platykurtic distribution shape). Notably, takeover speed alone is not a sufficient indicator of performance; hasty takeovers can reduce post-ToC control stability [27]. Interestingly, fatigue-related impacts on takeover performance appear inconclusive [57]. For example, [83] found that most re-engagement and braking metrics did not differ significantly between sleep-deprived and long-drive groups, although a subset of sleep-deprived drivers delayed automation deactivation. Consistent with this, takeover timing and control quality are shaped by situational complexity and driver factors: traffic density has a strong effect on prolonging takeover times and degrade post-ToC performance [26], whilst targeted training improves post-takeover control [55]. In addition, ToR lead time and trigger policy, HMI modality and content, scenario complexity, and environmental conditions (e.g., lighting, weather, work-zones) are consistently associated with both timing and post-ToC performances [86, 45, 52, 37].

2.2.2 System-related factors

ODD boundaries and ToR timing

From the system perspective, the handling of ODD exits governs when and how ToRs are issued. Two complementary classifications are useful here: an urgency-based view that distinguishes non-emergency transitions at system limits from emergencies caused by malfunctions or imminent hazards, and a cause-based view that separates ToRs triggered by inherent ODD limitations from those prompted by system faults [48]. Taken together, these perspectives frame the ODD boundary concept, which has been described as a "conditional hyperspace" [48]. In practice, standards such as PAS 1883 and ASAM OpenODD provide machine-readable taxonomies of ODD attributes, enabling precise specification of boundaries [8, 5]. Building on this, the scene–situation–scenario semantics of [77] and the functional–logical–concrete categorisation in [46] turn those boundaries into parameterised, executable scenarios. This makes boundary crossings explicit ToR triggers with controllable timing, i.e., the ToR can be issued at a specified lead time and takeover location, thereby scheduling the ToC. Subsequently, use-case–driven ODD treatments then embed these boundaries in the safety case to justify when ToRs are warranted [28, 88]. Examples of ODD boundaries include environmental limits (rain, fog, glare), dynamic infrastructure states (missing lane markings, roadworks), road-class or geometry constraints, and internal thresholds (e.g., perception performance, localisation confidence, map currency).

Scenario complexity and local traffic state

In addition to ODD limits, ToC demands are shaped by the driving scenario. Higher scenario complexity (e.g. number and type of agents, road geometry and environmental constraints) tends to worsen takeover performance [14]. A common construct in human-factors studies of automated driving is the *time budget* (TB), defined as the interval between a ToR and the moment a critical event or system limit would be reached (also termed "lead time" or "time buffer") [45, 14]. In practice, TB is often instantiated using time-to-boundary measures such as time to collision (TTC), and it serves as a proxy for urgency and thus for the required ToR lead time. Other traffic-state metrics (e.g. density, speed variance, time headway, lane-change frequency) can be evaluated at run time as complementary indicators, though their value is context-dependent. Empirically, a shorter TB (i.e. higher urgency) yields faster but poorer takeovers, whereas longer TBs improve situation awareness and control [14, 89, 38]. Meta-analytic evidence further shows that TB is the dominant driver of takeover time, while the presence of surrounding traffic increases mean

takeover times ($\approx +0.5\,\mathrm{s}$ on average) and higher automation (Level 3+) is associated with longer takeover times than partial automation (Level 2) [93]. Deceleration is the dominant initial response by human drivers after a ToR, especially in dense or complex traffic [19, 60, 7]. In operational terms, this implies that automation-side ToC planning should temporarily raise the minimum time headway above the nominal car-following target to accommodate longer driver response times, short TBs in complex scenarios, and the expected initial braking response by human drivers.

Together, these findings establish system-related and human-centred factors that shape ToC outcomes. Next, Section 2.3 reviews representative models covering driver-motivated and system-initiated transitions.

2.3 Existing Models of Control Transitions

Numerous modelling approaches from human factors, human sciences, engineering, and traffic research depict aspects of manual and automated driving with varying levels of detail, depending on the purpose and perspective of the investigation. Comprehensive overviews across these domains can be found, for example, in [32] and, more recently, in [35], both illustrating the complexity and diversity of modelling goals and frameworks, and tracing developments from the origins of driver modelling to present-day paradigms.

A large subset of models has emerged from the human factors domain, often focusing on driver-motivated takeover behaviour, for example predicting performance based on driver state, modelling decision-making processes under varying urgency, or simulating interaction strategies in shared-control scenarios. These typically emphasise aspects such as cognitive load, response times, gaze behaviour, or task-switching during transitions initiated by the driver. In parallel, other work has modelled system-initiated control transitions, either in response to critical situations or in planned, non-critical contexts. Table 2.1 provides a selection of representative models covering partial aspects of the control transition process.

While SAE Level 4 automation is designed to perform the complete dynamic driving task within its ODD, research has still explored situations in which control transitions may involve a human operator. In particular, the remote operation research field examines how human supervisors, typically located outside the vehicle, can intervene in non-emergency or degraded situations. For example, [59] use operational sequence diagrams to model interaction flows for four distinct levels of remote operation, from monitoring to full remote driving. Other studies propose control centre architectures and state diagrams to structure automation $\leftrightarrow$ remote operator handovers, or frameworks for selecting the appropriate remote operation mode based on ODD constraints. Although these models address control transitions,

Table 2.1: Selection of models on partial aspects of control transitions

Initiation	**Urgency**	**Level 1/2**	**Level 3**
Driver	Emergency	risk allostasis theory [81]; task demand and anticipation [10, 11]; active inference [87]; perceptual-gating action-decision model [44]	ACT-R + equation model [31]; perceptual-gating action-decision model [44]
	Non-Emergency	conceptual shared-control model of haptic steering [61]	QN-MHP + freeway-exit response model [74]
System	Emergency	state-transition and threshold-based driver response model for ACC failures [71]	regression-based ToC performance meta-model from pooled simulator data [25]
	Non-Emergency	task demand and anticipation [10, 11]	Machine learning takeover performance prediction model based on driving simulator studies [18]

they differ fundamentally from Level 3 scenarios in that the human is not expected to be a continuously engaged fallback driver, but rather an on-demand supervisor engaged only in specific, often rare, situations. Automation-centred research focuses on fallback strategies without human involvement, such as state-machine mappings of failures to predefined minimal-risk manoeuvres [92], sensor-failure fallback control to safe areas [90], or dynamic driving task fallback policies prioritising safe locations over immediate stops [20].

Taken together, the models surveyed above are primarily data-driven approaches (e.g., regression or machine-learning performance models) but rarely provide an operational, end-to-end state representation of the joint driver–automation system for Level 3 operation. There is no widely adopted state model for ToCs linked to traffic-flow analyses; most work treats pre- or post-takeover performance in isolation or specifies fallback logic for failed transitions. The next Section 2.4 therefore turns to modelling transitions in the context of traffic simulation.

2.4 Modelling Transitions in Traffic Simulation

As reviewed, most human-centred modelling efforts focus on the driver's behaviour after a takeover request, including perception, decision-making, and subsequent control execution, all from a human factors perspective. In contrast, research from the automation system perspective, often with a strong technical focus, is primarily concerned with the ability of the automation to recognise constraints or failures,

issue takeover requests in a timely manner, and manage failures through minimal risk manoeuvres. Although both perspectives are important, they remain largely disconnected. There is no established modelling framework that covers the entire transition process on both the human and automation sides in a way that allows impacts at the traffic level to be studied. From a traffic sciences perspective, this represents a clear modelling gap which the microscopic traffic simulator SUMO [4] addresses with its ToC model, presented by [42], enabling the representation of transition behaviour in the context of traffic flow and the investigation of its effects on safety, capacity, and efficiency.

Control transitions in SUMO are implemented through a dedicated vehicle device, which manages the switching between control regimes (manual $\leftrightarrow$ automated driving) by altering parameter sets in the underlying car-following or lane-changing models, or by substituting one model for another. The device also represents processes during a control transition, such as the automated preparation for a takeover and the temporary decrease in human driving performance after manual control is resumed. The entire model is essentially a state machine model, as illustrated in Figure 2.2.

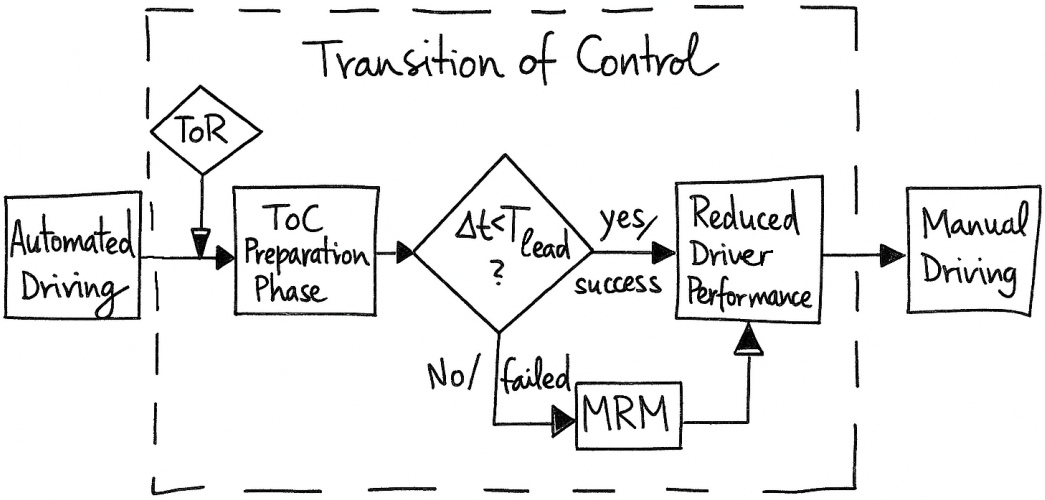

Figure 2.2: State machine of the Transition of Control in SUMO. Sequence from `Automated Driving` via `ToR` to `ToC Preparation Phase`. Decision node: available time Δt compared to threshold T_{lead}. Successful transition: `Reduced Driver Performance` followed by `Manual Driving`. Failed transition: activation of an `MRM` (arrow depicts the driver-intervention case; otherwise the MRM persists to a minimum-risk condition, typically standstill; not shown).

The *ToC Preparation Phase* is realised by a gap control mechanism, which can increase the desired time gap to a leading vehicle. It should be noted that such a mechanism is not mandatory in Level 2, which requires the human driver to be in a permanent, instantaneous state of readiness. The current regulations for Level 3 do

not yet specify in detail how exactly a vehicle automation system should establish a safe vehicle state for the takeover during the transition phase. [1] Depending on the parameterisation of the dedicated gap controller, a ToC can be modelled as Level 1/2 or as Level 3, with the latter entailing an increased headway, as illustrated schematically in Figure 2.3.

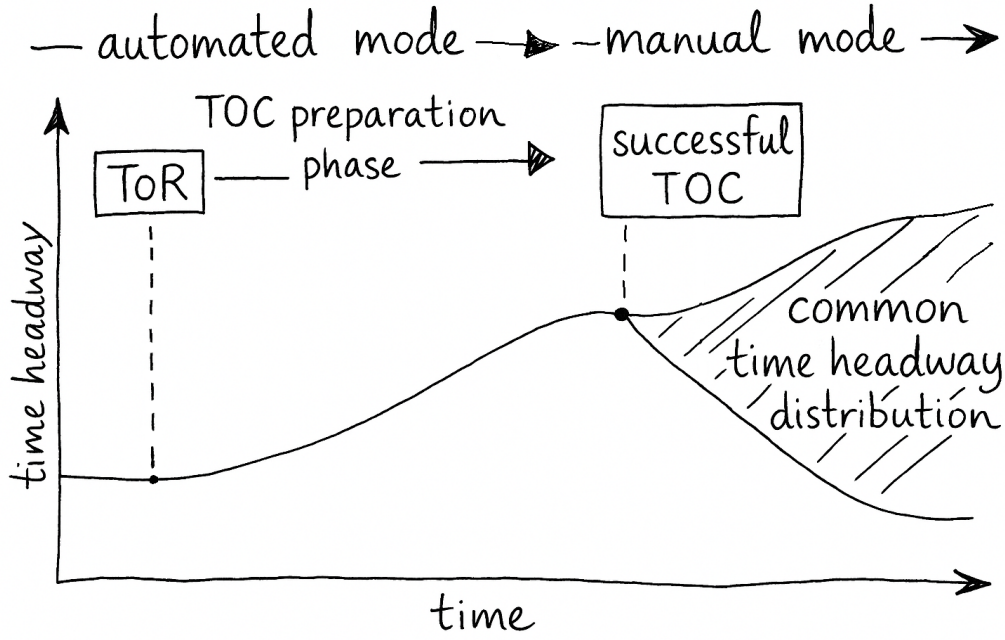

Figure 2.3: Time headway increase during the ToC preparation phase in a car-following situation, followed by convergence towards the common distribution of manually driven headways after takeover.

The state of reduced driving performance after a downward transition is modelled through a driver state mechanism that accounts for time-variant driver awareness, with linear recovery from an initially lower level to full performance after takeover. In case of a failed transition, the system initiates an MRM toward a minimum-risk condition. If the driver resumes control before the MRM's terminal state, this realised outcome is a short MRM and the process continues with *Reduced Driver Performance*. Otherwise, the MRM persists to standstill at a predefined constant deceleration rate.[2]

Bringing these mechanisms together, the SUMO ToC device provides an executable, traffic-level representation of Level 3 control transitions, integrating system-

[1] The Regulation No. 157 for Level 3 automated lane keeping systems (ALKS) by the UNECE [78] e.g. only states in their section 5.4.3. "During the transition phase the system shall continue to operate. The system **may reduce the speed of the vehicle** to ensure its safe operation but shall not bring it to standstill unless required by the situation."

[2] A full technical description of the configuration and parameterisation of SUMO's ToC device can be found in the online documentation: https://sumo.dlr.de/docs/ToC_Device.html.

initiated ToRs, preparation behaviour, temporary post-ToC performance, and the MRM fallback within a single state machine framework. The timing of ToRs can also be controlled programmatically via SUMO's TraCI interface. It supports per-vehicle parameterisation and controlled variation of preparation policies (e.g., headway increase), trigger timing, and driver-state recovery, thereby enabling experiments with mixed-autonomy traffic compositions.

2.5 Traffic-Level Effects of Level 3 Control Transitions

ToC-related research faces three constraints: (i) it predominantly addresses human-centred aspects (cf. Section 2.2); (ii) incident-driven traffic safety studies [33, 64, 73, 15, 1] rely on disengagement reports (e.g., California DMV [9]) and crash reporting under NHTSA's Standing General Order [53], which together largely cover Level 2 or Level 4 incidents and higher-level testing rather than Level 3 operations; and (iii) real-world Level 3 exposure remains limited because systems have only recently been introduced and are rarely available, as reflected in geographically localised on-road testing in the United States [54]. Comparable public reporting of disengagements or automation-level crash data is largely absent in Europe and Japan; the European Transport Safety Council notes that no mandatory public reporting exists [23]. As a result, publicly available data on routine Level 3 operations are limited, in particular for system-initiated, non-emergency ToCs at scale. Consequently, the traffic-level effects of these transitions remain under-investigated.

Given these constraints, this thesis adopts a simulation-based approach to quantify such effects. Under the current version of Regulation No. 157 for Level 3 automated lane keeping systems (ALKS) by the UNECE [78], the Level 3 ODD is confined to motorways. Scenario-based developments on rural and urban roads have only recently been explored in limited form by projects such as L3Pilot [36] and Hi-Drive [17]. Accordingly, Chapter 3 presents simulation studies to assess network-level effects of control transitions. In practice, such transitions occur in mixed fleets with only partial Level 3 penetration; therefore, the simulations vary Level 3 penetration rates across scenarios. To isolate the effects of control transitions, other behavioural or dynamic factors specific to Level 3 (usage rates, driver-type heterogeneity, road type, infrastructure compatibility, ODD readiness) are either held constant or embedded within the simulation use-case definitions. Full driver compliance is assumed. The overall objective is to quantify mitigable yet unavoidable ToC-related traffic impacts at the network level.

At the traffic-flow scale, explicit treatments of ToC effects are scarce; a brief review

of the few available studies clarifies mechanisms relevant for the simulation design. An analytic study by [85] derives a string-stability condition for mixed traffic with takeover requests and validates it on a ring road, showing that increasing takeover incidence near deceleration disturbances (e.g., drops in the lead vehicle's speed) degrades stability, with TTC thresholds and Level 3 vehicle penetration governing the effect size. Complementing this, [39] simulate mixed traffic at a motorway merge and report that moderate takeover times (5–7 s) support stable flow, whereas very short or very long takeover times (2 s or 10 s) destabilise it. Higher penetration of conditionally or fully automated vehicles (AV) expands the stability region. Whereas these studies emphasise post-takeover dynamics, the present thesis concentrates on the ToC preparation phase in dense traffic.

Within this preparation-phase scope, a limiting factor in exploring ToC effects on traffic flow is failed control transitions that result in extended MRMs. Such fallback manoeuvres behave like incidents when they culminate in a standstill within a running lane and can dominate traffic dynamics. Unless the vehicle is guided to the hard shoulder, such MRMs trigger queue formation that obscures the effects of otherwise predominant successful ToCs. These unsuccessful ToC events are comparatively rare and vehicle-specific, requiring targeted vehicle and infrastructure interventions. Accordingly, this thesis focuses on preparation-phase effects. MRM impacts and traffic management measures to mitigate them are treated elsewhere, for example in use-case-based simulations [43, 3]. To isolate potentially mitigable, cumulative effects of system-initiated, non-emergency Level 3 transitions, the simulations assume mostly successful takeovers and, in a small fraction of cases, only short MRMs with human response times ≤ 3 s that do not result in complete lane blockages.

Methodologically, the simulation studies presented in Chapter 3 adopt a SUMO-based approach with per-vehicle parameterisation (vehicle type and ToC behaviour) and parameter values sampled from distributions to represent heterogeneous traffic. Most importantly, consistent with Regulation No. 157, Level 3 automated vehicles are assigned a desired time headway $\tau_{\mathrm{L3}} = 1.6$ s, which is deliberately higher than values often assumed for automated vehicles in the literature (e.g., sub-second cooperative adaptive cruise control (CACC)–type settings or Level 4 AVs) and typically above reported human desired headways on motorways [84]. These settings operationalise the regulatory constraints in the experimental design across all studies.

Chapter synthesis

This chapter defined the core concepts (ToC types, ODD and ToRs), surveyed human- and system-centred determinants of takeover performance, reviewed representative aspects of ToC modelling, and motivated an executable framework for traffic-level analysis using SUMO's ToC device. Building on this foundation, Chapter 3 distils modelling approaches and key findings from the associated publications and presents the three main papers quantifying traffic-level impacts on efficiency, safety and capacity. Chapter 4 synthesises overarching issues and complements these results with additional analyses and follow-up experiments. In doing so, the thesis proceeds to address RQ 1 (efficiency), RQ 2 (safety) and RQ 3 (capacity) in turn.

3 Publications

The research activities on ToCs began within the framework of the EC-funded project TransAID [76]. Several publications and technical reports emerged from this work and serve as a basis for the main research contributions presented in the subsequent sections. Initial implementations, modelling approaches, simulation setups, parameterisation schemes and key findings from TransAID laid the groundwork for the developments and analyses discussed in the three main papers.

3.1 Associated Publications

3.1.1 From Automated to Manual - Modelling Control Transitions with SUMO

The following publication marks a starting point for the research presented in this thesis, as it consolidates the core outcomes from the initial phase of the TransAID project:

1. Leonhard Lücken, Evangelos Mintsis, Kallirroi Porfyri, Robert Alms, Yun-Pang Flötteröd, and Dimitrios Koutras. "From Automated to Manual - Modeling Control Transitions with SUMO". in: *SUMO User Conference 2019*. Ed. by Melanie Weber, Laura Bieker-Walz, Robert Hilbrich, and Michael Behrisch. Vol. 62. EPiC Series in Computing. EasyChair, 2019, pp. 124–144. DOI: 10.29007/sfgk

This paper [42] presents the modelling framework for simulating ToCs in SUMO as outlined in Section 2.4. Moreover, two traffic management scenarios are introduced to demonstrate the impact of ToCs and the effectiveness of mitigation strategies based on vehicle-to-everything (V2X) messages that roadside units (RSU) can transmit to connected automated vehicles (CAV). In the first, providing CAVs with path information via RSUs reduces unnecessary ToCs near obstacles, improving safety without compromising throughput or emissions. In the second use case, labelled 5.1, CAVs in a two-lane motorway scenario approach an area where their ODD ends, a no-automated-driving zone (No-AD zone), cf. Figure 3.1.

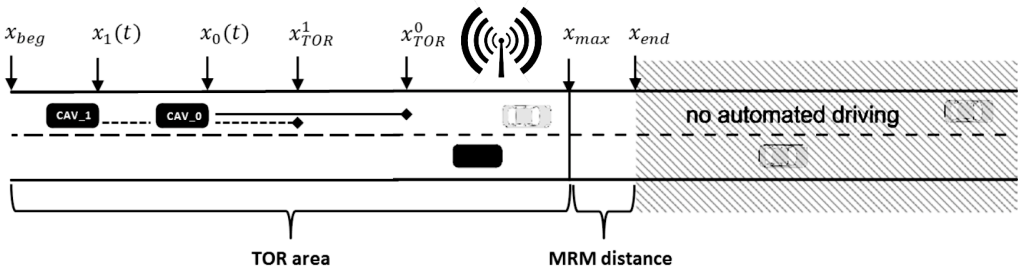

Figure 3.1: Illustration of use case 5.1 from Figure 5 in [42]: CAVs receive a system-initiated ToC at position x_{ToR} by a RSU due to an approaching area where automated driving is not permitted. Drivers must resume control before reaching $x_{\max}$ to avoid triggering a MRM within the remaining distance up to x_{end}.

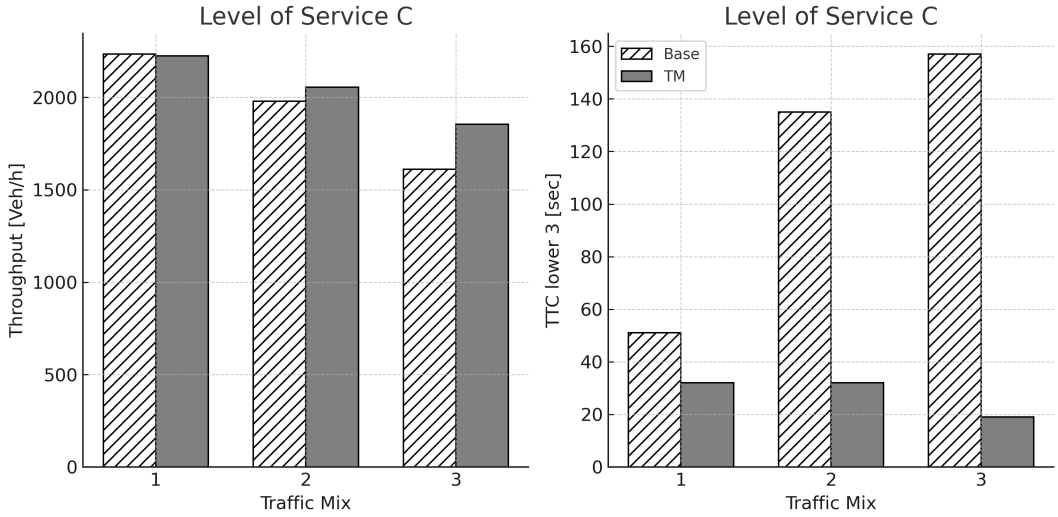

Figure 3.2: Excerpt of results from Figure 8 in [42]: Throughput and critical time headways (TTC $< 3\,\mathrm{s}$) for different traffic mixes under Level of Service C. The baseline scenario (Base) is shown with hatching, the managed case (TM) in gray.

A heuristic ToR scheduling algorithm distributes the timing of ToCs upstream of the No-AD zone. Executed at the RSU-level, the algorithm processes the observed longitudinal states of approaching vehicles, the leader speed, and a local density factor, and returns individual ToR times t_i and positions x_i^{ToR}. Driver response and takeover execution are handled by the SUMO ToC model (cf. Section 2.4). As detailed in [43], Algorithm 1 [3] implements three steps: (1) compute a speed-dependent interval Δt from the difference between the target ToR gap and the automated-driving gap together with a braking rate b_{MRM} during an eventual MRM; (2) choose a ToR point for the leading vehicle within $[x_{\mathrm{beg}}, x_{\max}]$ as a function of the density factor and assign followers at $t_i = t_0 - i\,\Delta t$, with positions for $i > 0$ computed as $x_i(t) + t_i(t)\,v_0(t)$; and (3) proportionally rescale all trigger positions with any change in $x_{\max}(t)$ so that all triggers remain within the managed segment. The effectiveness of this strategy is

[3]The algorithm was implemented in Python and interfaces with SUMO via the TraCI API.

illustrated in Figure 3.2, showing improved throughput and reduced TTC numbers compared to the unmanaged baseline.

More importantly, the results from this scenario uncover a basic phenomenon: quasi-synchronous ToCs at a specific location, e.g. at the end of a local ODD as in these experiments, lead to bottleneck-like congestion under high traffic demand. If not managed properly, as also shown in this paper, such ToCs reduce capacity. This observation motivates the analyses in Sections 3.2–3.4. Across these studies, Algorithm 1 is used as the reference traffic-management baseline to isolate traffic-level effects of different ToC patterns and to evaluate mitigation strategies.

Algorithm 1: ToR scheduling algorithm based on [43, Section 3.1.5.2.1]

Parameters : $\mathrm{SPACING}_{\mathrm{ToR}}$, $\mathrm{TIMEGAP}_{\mathrm{ToR}}$, $\mathrm{SPACING}_A$, $\mathrm{TIMEGAP}_A$, b_{MRM}, x_{beg}
`// fixed parameters from the section, Table 12//`

Inputs : current states $x_i(t)$, $v_i(t)$ for vehicles $i = 0, \ldots, n$ (front to back);
leader index 0 with speed $v_0(t)$; density factor $\rho(t) \in [0, 1]$;
upstream limit $x_{\max}(t)$.

Outputs : updated ToR times $t_i(t)$ and ToR positions $x_i^{\mathrm{ToR}}(t)$.

while *simulation running* **do**

`// Step 1: compute the current ToR interval` $\Delta t(t)$

$g_1(t) \leftarrow \mathrm{SPACING}_{\mathrm{ToR}} + \mathrm{TIMEGAP}_{\mathrm{ToR}} \cdot v_0(t)$ `// Eq. (4)`

$g_A(t) \leftarrow \mathrm{SPACING}_A + \mathrm{TIMEGAP}_A \cdot v_0(t)$ `// Eq. (5)`

$D_g(t) \leftarrow g_1(t) - g_A(t)$ `// Eq. (3)`

$\Delta t(t) \leftarrow \sqrt{2\, D_g(t)\, b_{\mathrm{MRM}}}$ `// Eq. (7)`

`// Step 2: schedule current ToR times and positions`

choose or update leader ToR time $t_0(t)$

$x_0^{\mathrm{ToR}}(t) \leftarrow \min\{\, x_{\max}(t),\, x_{\mathrm{beg}} + \rho(t)\,[x_{\max}(t) - x_{\mathrm{beg}}]\,\}$ `// Eq. (9)`

for $i \leftarrow 0$ **to** n **do**

$t_i(t) \leftarrow t_0(t) - i \cdot \Delta t(t)$ `// Eq. (8)`

if $i > 0$ **then**

$x_i^{\mathrm{ToR}}(t) \leftarrow x_i(t) + t_i(t) \cdot v_0(t)$ `// Eq. (10)`

`// Step 3: dynamic scaling with moving` $x_{\max}(t)$

for $i \leftarrow 0$ **to** n **do**

$\theta_i(t) \leftarrow \dfrac{x_i^{\mathrm{ToR}}(t) - x_i(t)}{x_{\max}(t) - x_i(t)}$ `// Eq. (11)`

$x_i^{\mathrm{ToR}}(t) \leftarrow x_i(t) + \theta_i(t)\,[x_{\max}(t) - x_i(t)]$ `// Eq. (12)`

`// issue a ToR when vehicle` i `reaches` $x_i^{\mathrm{ToR}}(t)$

3.1.2 Related Publications

The following reports and papers supply the technical basis for the SUMO studies in this thesis. They define the automation and driver-state models, the traffic/vehicle parameterisation and simulation setup, and provide context to the RSU-centred V2X mechanisms adopted in the motorway use cases:

2. Evangelos Mintsis, Dimitris Koutras, Kallirroi Porfyri, Evangelos Mitsakis, Leonhard Lücken, Jakob Erdmann, Yun-Pang Flötteröd, Robert Alms, Michele Rondinone, Sven Maerivoet, Kristof Carlier, Xiaoyun Zhang, Robbin Blokpoel, Martijn Harmenzon, and Steven Boerma. *TransAID Deliverable 3.1 - Modelling, simulation and assessment of vehicle automations and automated vehicles' driver behaviour in mixed traffic.* Tech. rep. European Commission, Sept. 2020. DOI: 10.3030/723390. URL: https://cordis.europa.eu/project/id/723390/results

    - Report no. 2 of [49] provides a detailed account of the modelling aspects of vehicle automation, including the implementation of adaptive cruise control (ACC) and cooperative adaptive cruise control (CACC) models [65], as well as the introduction of the driver state model for ToCs in SUMO. The parameterisation schemes for traffic mixes and vehicle type definitions (light goods vehicles (LGV); heavy goods vehicles (HGV); legacy/manual vehicles (MV); CAV) described in this report were partially adopted for the simulation experiments described in Sections 3.2–3.4.

3. Sven Maerivoet, Lars Akkermans, Kristof Carlier, Péter Pápics, Bart Ons, Stef Tourwé, Robert Alms, Yun-Pang Flötteröd, Leonhard Lücken, Evangelos Mintsis, Vasilios Karagounis, Dimitrios Koutras, Anton Wijbenga, Jaap Vreeswijk, Alejandro Correa, Xiaoyun Zhang, and Robbin Blokpoel. *TransAID Deliverable 4.2 – Preliminary Simulation and Assessment of Enhanced Traffic Management Measures.* Tech. rep. European Commission, Oct. 2020. DOI: 10.3030/723390. URL: https://cordis.europa.eu/project/id/723390/results

    - Report no. 3 of [43] builds upon no. 2 by presenting an extensive simulation setup, adopting models and varied vehicle and traffic parameter schemes from no. 2, forming the basis for a simulation study comprising ten use cases. The results and findings of use case 5.1 in particular, as described in Section 3.1.1, have been picked up in the research of Section 3.2.

4. Alejandro Correa, Robert Alms, Javier Gozalvez, Miguel Sepulcre, Michele Rondinone, Robbin Blokpoel, Leonhard Lücken, and Gokulnath Thandavarayan. "Infrastructure Support for Cooperative Maneuvers in Connected and Automated Driving". In: *2019 IEEE Intelligent Vehicles Symposium (IV)*. 2019, pp. 20–25. DOI: 10.1109/IVS.2019.8814044

    - Publication no. 4 offers a condensed evaluation of a then-novel V2X message type, the Manoeuvre Coordination Message (MCM) [22], that conceptually and technically enables traffic management strategies to coordinate CAVs via RSUs in critical cases, exemplified here with use case 5.1.

Outside the TransAID context, an additional publication [2] examines safety-related aspects for Level 3-compliant ALKS that fall under UNECE Regulation No. 157 [78].

5. Robert Alms, Benjamin Couéraud, and Peter Wagner. "Perspectives on an ALKS Model in SUMO". in: *SUMO Conference Proceedings* 5 (July 2024), pp. 269–285. DOI: 10.52825/scp.v5i.1198

    - The paper no. 5 does not discuss ToCs in particular, since they are not considered in the safety assessment of the regulation, but it touches on the modelling challenges regarding Level 3-specific vehicle behaviour in SUMO, i.e., emulating a lane-change detection mechanism resembling the sensor-based model described in Regulation 157 which is critical to the premises of the papers presented in Sections 3.3 and 3.4.

Taken together, the associated publications establish the modelling and simulation basis used throughout this thesis: the SUMO framework for system-initiated control transitions, the RSU-side ToR scheduling heuristic, and the parameterisation schemes for vehicles, traffic mixes and V2X coordination. Use case 5.1 provides the empirical baseline on which subsequent analyses build. Building on this foundation, Section 3.2 presents the first study, which examines effects on traffic efficiency and compares ToR-scheduling strategies for mitigating ToC-induced impacts.

3.2 Paper I - Reinforcement Learning-Based Traffic Control: Mitigating the Adverse Impacts of Control Transitions

Originally published as:

Scope and Contribution of Paper I

This paper extends the preliminary work of [42] by focusing on use case 5.1, in which quasi-synchronous ToCs occur at the boundary of a No-AD zone. Based on the previously established unmanaged versus managed cases, several coordinated traffic management strategies are compared. The optimisation problem is formulated as a reinforcement learning (RL) task, framed as a trade-off between maintaining traffic efficiency and prolonging automated driving. In RL, an agent learns a policy through repeated interaction with the environment, guided by a reward signal. The policy denotes the mapping from states to actions, complemented in actor–critic methods by a critic estimating the expected return of those actions. Different algorithms implement this interaction in distinct ways: Proximal Policy Optimisation (PPO) [69] is an on-policy gradient method that requires fresh data for each update, whereas Twin Delayed Deep Deterministic Policy Gradient (TD3) [24] is an off-policy actor–critic method that reuses past experience via a replay buffer and employs twin critics for stability. In constructing the simulation environment, traffic mixes and parameterisation schemes were adopted from [49], and the motorway was discretised into cells (each $\approx 350\,\mathrm{m}$ long) containing local traffic flow information, which define the observation space of a Markov Decision Process (MDP). Two alternative reward functions were designed, and the action space was defined by a transmission rate of ToRs $a \in [0, 1]$.

Initial training experiments prior to the paper submission, motivated by the rapid improvement of the PPO algorithm in the machine learning domain at the time,[4]

[4]From DQN's 2015 "human-level" Atari results [50] to PPO's large-scale success at OpenAI in 2019 [56], and further evidence of PPO's effectiveness in cooperative multi-agent games [91], similar

failed due to insufficient parallelisation and the poor performance of PPO in this RL setup using Stable Baselines 2 (SB2) [29][5]. Therefore, with a revised formalisation of the MDP, the paper reports experiments with the sample-efficient TD3 algorithm. In this case, the continuous action formulation meant that the transmission rate was modelled as a real-valued control variable rather than a discrete choice. Training was conducted on the high-demand scenario LoD C with traffic mix 3, representing an AV share of 40%. Aggregated results for travel time and average CAV distance showed that RL models 1 and 2 achieved a balanced compromise between traffic efficiency and automated driving continuation in the training scenario. This led to the overall conclusion that RL-based controllers are at least on par with traditional heuristic traffic management approaches.

With respect to RQ 1, unmanaged quasi-synchronous ToCs in dense traffic (LoD C, Mix 3) cause a clear drop in traffic efficiency. Heuristic staggering reduces this loss relative to the unmanaged baseline. The RL-based ToR-scheduling policy delivers comparable gains in mean travel time and better preserves automated driving distance. In this setup, ToCs materially degrade efficiency when left unmanaged, and ToR scheduling is an effective traffic management measure. Further evidence and caveats are discussed in Section 4.1.

RL techniques have been adopted in the traffic domain for traffic control and mixed-autonomy traffic management (e.g., [82, 80, 12, 6]).

[5]The PPO algorithm and an up-to-date version of the Stable Baselines library (SB3) is discussed later in Section 4.1.3.

IEEE Open Journal of
Intelligent Transportation
Systems

Received 18 August 2021; revised 10 February 2022; accepted 7 March 2022. Date of publication 11 March 2022; date of current version 22 March 2022.

Digital Object Identifier 10.1109/OJITS.2022.3158688

Reinforcement Learning-Based Traffic Control: Mitigating the Adverse Impacts of Control Transitions

ROBERT ALMS[1], ARISTEIDIS NOULIS[2], EVANGELOS MINTSIS[3], LEONHARD LÜCKEN[4], AND PETER WAGNER[1,5]

[1]Institute of Transportation Systems, German Aerospace Center (DLR), 12489 Berlin, Germany

[2]ARRC, Technology Innovation Institute (TII), Abu Dhabi, UAE

[3]Hellenic Institute of Transport, Centre for Research and Technology Hellas, 57001 Thessaloniki, Greece

[4]ICBM, University of Oldenburg, 26111 Oldenburg, Germany

[5]Institute of Land and Sea Transport Systems, TU Berlin, 10587 Berlin, Germany

CORRESPONDING AUTHOR: R. ALMS (e-mail: robert.alms@dlr.de)

This work was supported in part by the EC Project TransAID under Grant 723390.

ABSTRACT An important aspect of automated driving is to handle situations where it fails or is not allowed in specific traffic situations. This case study explores means, by which control transitions in a mixed autonomy system can be organized in order to minimize their adverse impact on traffic flow. We assess a number of different approaches for a coordinated management of transitions, covering classic traffic management paradigms and AI-driven controls. We demonstrate that they yield excellent results when compared to a do-nothing scenario. This text further details a model for control transitions that is the basis for the simulation study presented. The results encourage the deployment of reinforcement learning on the control problem for a scenario with mandatory take-over requests.

INDEX TERMS Connected automated vehicles (CAV), reinforcement learning (RL), take-over request (ToR), traffic management (TM), transition of control (ToC).

NOMENCLATURE

AV	Automated vehicle
CAV	Connected automated vehicle
CV	Connected vehicle
LoD	Level of demand
MRM	Minimum risk manoeuvre
MV	Manual vehicle
MDP	Markov decision process
No-AD	No automated driving
RL	Reinforcement learning
RSI	Roadside infrastructure
TM	Traffic management
TMC	Traffic management center
ToC	Transition of control
ToR	Take-over request.

The review of this article was arranged by Associate Editor Jia Hu.

I. INTRODUCTION

THE TREND towards vehicle automation and connectivity between vehicles (V2V) or infrastructure (V2I) implies a need for traffic management approaches to deal with emerging complications in future mixed traffic situations. In such scenarios, connected automated vehicles (CAV) can be addressed individually by V2I technology, which differs considerably from classic traffic management tasks that organize large numbers of road users with standard control measures like, e.g., signal control or ramp metering. Based on this communication, CAVs hold the potential to pose as sensors and actuators within the traffic system simultaneously. This opens up prospects for novel traffic management schemes.

Moreover, with the progressing deployment of CAVs that provide state-of-the-art level-two functionalities (see SAE taxonomy for automated vehicles [1]), the traffic system

will gradually turn into a system of mixed autonomy. Such a system presents various challenges in terms of traffic efficiency and safety when human drivers and partly to fully automated vehicles (AV) share the same road space. In particular, automation disengagements are of concern, i.e., when a human driver has to operate as a fallback for a failed vehicle automation and needs to respond in a proper and timely manner to take back the driving task [2]. This safety critical process, a so-called downward *transition of control* (ToC) [3], is an increasingly important studied research topic, especially from the perspective of manufacturers on how to design respective takeover strategies in highly automated vehicles [4], [5], [6].

In contrast, the macroscopic effects on overall traffic, which even successful downward ToCs may induce when occurring frequently in certain traffic situations and areas are less investigated so far. Therefore, in this paper we consider the issue of such transition areas from a traffic management perspective with a specific focus on:

1) how to model, simulate and manage downward ToCs, detailing some of the related work of the EC project *TransAID*, and on,
2) how to design a traffic management control that mitigates the adverse effects of downward ToCs on traffic with the help of tools from artificial intelligence compared to a more traditional approach.

To the best of the authors' knowledge, there is only one other publication presenting a model for ToCs that conducted simulations on traffic performance [7], but with a rather limited scope on restricting the lane change behaviour. Thus, we present a case study based on a novel ToC model, that demonstrates possible outcomes in future scenarios of mixed autonomy when downward ToCs contribute detrimental to the traffic performance on a macroscopic level. Our study introduces several approaches on managing those adverse impacts.

We apply *reinforcement learning* (RL) which is a concept in machine learning that formalizes a control task in form of a *Markov Decision Process* (MDP) [8] to maximize a reward in a trial-and-error learning process. Reference [9] points out that the terminus RL is a class of solution methods in machine learning as well as a research field of these solutions that work well on a problem. In that sense, this work simply uses RL as a method to solve a control task for traffic management rather than researching the problem of RL itself. RL-based methods have been previously adopted to optimize traffic light performance [10], control ramp meters [11], enable eco-driving along signalized corridors [12], deliver personalized driving policies [13], or facilitate big-data driven intelligent traffic management [14]. However, the RL-based management of downward ToCs upstream of a no-automated-driving zone (No-AD zone) has not been targeted by RL approaches up to date.

The rest of the paper is organized as follows. In Section II we introduce the modeling of ToCs and discuss the outcome of respective simulation results. Further, we briefly review some related publications in the context of traffic management of CAVs with the application of RL. Section III describes the computer experimental setup. That is, we define the traffic control task and formulate the MDP for the RL experiment. Section IV presents the simulation results and provides a discussion of the obtained results. Finally, in Section V we present our conclusions from this study and point to future research directions.

II. RELATED WORK

A. TRANSITIONS OF CONTROL

ToCs constitute overarching processes that govern bi-directional shifts of authority between the driver and the AV. In case of downward ToCs, factors endogenous or exogenous to the AV may force vehicle automation to disengage and request driver's intervention for resuming AV's control. The signal (audio, visual, haptic or combination of the latter) from the vehicle automation side that notifies the driver for the need to re-engage in the primary driving tasks is defined as a take-over request (ToR). A successful downward ToC is completed as soon as the driver has re-engaged and his/her situational awareness and driving skills are fully restored. If the downward ToC is unsuccessful, namely the driver does not respond to ToR within the available lead time, the AV stops as safely as possible via a minimum risk manoeuvre (MRM). In case of upward ToCs, the driver hands over control to AV within its Operational Design Domain (ODD) via the activation of its automated driving systems. For example, the manufacturer Daimler recently announced that it will introduce a conditionally automated level-3 system in 2022, which deploys this depicted takeover strategy [15].

Failure from the driver's side to react in a timely manner to ToRs has been linked with fatal crashes in reports from both the U.S. National Highway Traffic Safety Administration (NHTSA) and the National Transportation Safety Board (NTSB) [16]. Given the adverse impacts of downward ToCs on safety, multiple studies have ventured to identify contributing factors to automated vehicle disengagements, by harnessing data collected for the disengagement and AV collision reports of the California Department of Motor Vehicles [17]–[19]. Findings from the latter studies indicate that vehicle-initiated disengagements are positively correlated with sensing and planning issues on the vehicle side and occur with increased frequency in high speed driving conditions or when behavior from other traffic participants becomes unpredictable and irregular.

Another research branch on downward ToCs has placed emphasis on determining human and AV system factors that affect driving performance during the ToC-preparation and post-ToC phases. References [20], [21] conducted literature review studies to identify factors that influence response time to ToRs (available lead time, involvement in secondary tasks, take-over request functionality etc.) and post-takeover vehicle control (braking-steering) in different traffic situations. Moreover, they reviewed existing models suitable for capturing the aforementioned artefacts of driver behavior prevailing

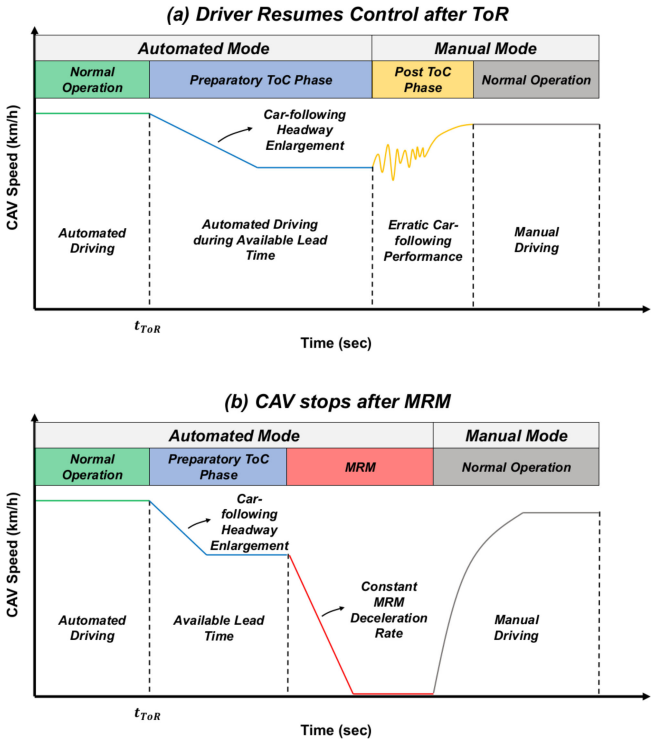

FIGURE 1. Illustration of the proposed ToC model for both cases: panel (a) shows a successful control transition; panel (b) presents an unsuccessful control transition resulting in a MRM.

in the course of downward ToCs. High fidelity driver models that can comprehensively capture behavioural processes during downward ToCs are essential for exhaustively studying their impacts and assessing possible mitigation measures via computer simulations. To this end, [22] proposed a simulation framework that incorporated human factors (task demand and capacity, situational awareness) in microscopic traffic models which enabled the investigation of complex driver-vehicle interactions occurring during downward ToCs [23]. To capture the symptomatic side of potential disruptions of smooth vehicle operation induced by downward ToCs while, at the same time, enabling large scale simulations, [24] developed a simplified, computationally efficient model for ToCs, which allows capturing statistical characteristics of the take-over performance of AV drivers.

The modelling and large scale simulation of planned downward ToCs constitute focal elements in the context of this study. Since our primary interest lies on the statistical distributions of the downward ToC characteristics and associated potential disruptions of the smooth traffic flow, rather than on the detailed psycho-physical processes underlying these, we employ the ToC model adopted from [24]. A generic description of the ToC model is provided below, while its detailed mathematical formulations can be found in [24].

The proposed ToC model as illustrated in Fig. 1 is based on a state machine that enables transitions between automated and manual driving modes. ToRs issued during normal operation of automated mode prompt the commencement of a preparatory ToC phase when automated driving can be explicitly supported for a confined time interval (available lead time). Upon expiration of the available lead time, two distinct outcomes are possible according to driver's response to ToR. Either the driver reacts to the ToR and resumes vehicle control (Fig. 1, panel (a)), or the AV is forced to enter a minimum risk condition and stop as safely as possible (Fig. 1, panel (b)). In the context of the state machine, the first case pertains to the transition from the preparatory to the post-ToC phase, where the driver may exhibit a reduced driving performance until she/he fully restores her/his driving skills (normal operation in manual mode). The second case pertains to the execution of a minimum risk manoeuvre that safely stops the AV.

Our modelling approach encompasses the enforcement of lane change abstinence, acceleration abstinence, and the establishment of enlarged and secure car-following headways via a gap control mechanism throughout the preparatory ToC phase for safety reasons. The augmentation process of car-following headways can be manipulated with the adaptation of several calibration parameters of the latter mechanism (headway change rate, maximum allowed deceleration, duration of gap opening manoeuvre) to attain the desired car-following behavior (new desired headway, duration the new desired headway is maintained). Driving performance during the post-ToC phase is determined based on a driver state model that embeds perception errors in the default car-following behavior of the microscopic traffic simulator SUMO [25]. Each driver is randomly assigned an initial situational awareness state when she/he enters the post-ToC phase and a situational awareness recovery rate that regulates the restoration of situational awareness until normal operation in manual mode is achieved. Erratic car-following behavior is triggered during the post-ToC phase according to perception specific action points designed for imperfect driving [26]–[29]. Moreover, the considered ToC model assumes a constant deceleration rate during the MRM which can either take place in the vehicle's current lane or stop the AV on the right-most lane via lane changes (if surrounding traffic conditions permit).

Finally, the ToC model presumes two different ways for issuing ToRs in SUMO. In the first case, the location of ToRs can be a priori specified via a dedicated SUMO functionality. In the second case, ToRs are issued dynamically according to AV planned manouvres and surrounding traffic conditions. In specific, if an AV encounters a dead-end lane and is forced to execute a lane change for strategic reasons while nearby vehicles block the AV intended manoeuvre, then the AV will dynamically issue a ToR. On this occasion, the location of dynamic ToR becomes a function of AV speed and distance to the dead-end. Furthermore, dynamic downward ToCs encompass dynamical sampling of driver response time which entails probabilistic estimation of MRM frequency as well.

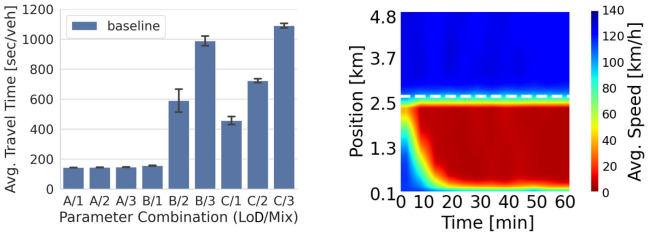

(a) Travel times for LoD A-C (b) Spatiotemporal speed diagram

FIGURE 2. **Baseline results of the TransAID use case study. Panel (a): Average travel times for different parameter combinations LoD/Mix, panel (b): Exemplary space-time-diagram for the mean speed of a single simulation run with LoD C and Mix 3. The white dashed line indicates the point from which all vehicles are obliged to drive manually.**

This approach for modelling ToCs makes it straightforward to embed the associated processes into microscopic traffic simulations of a broad variety of traffic scenarios at a high computational efficiency, such that the macroscopic assessment of impacts on traffic operations via microscopic simulation software can be achieved. A simulation analysis encompassing mixed fleet scenarios indicated that downward ToCs can induce adverse impacts on traffic efficiency, conflict risk and the environment in a variety of traffic situations (lane closure, road works, highway merge/diverge sections, no automated driving zones) [30]. Key findings of this study on the impact of downward ToCs on traffic flow upstream of a No-AD zone in the absence of any vehicle specific managing intervention from a Traffic Management Center (TMC)[1] are displayed in Fig. 2. In Panel (a) the average travel time of a single vehicle required to pass through the whole simulated road segment (cf. Fig. 3) is reported for an array of scenarios. These scenarios differ in the assumed level of demand (LoD), and the composition of the traffic, i.e., the percentage of automated vehicles. The LoD was varied ranging over the categories A, B, and C, and the percentage of different types of automated vehicles varied over Mix 1 (30%), Mix 2 (50%), and Mix 3 (80%), see Section III-C for details. Each scenario is assigned an ID composed of the corresponding LoD and Mix code, i.e., scenario C2 corresponds to LoD C and Mix 2. It can be observed that the do-nothing scenarios exhibit severe traffic jams beyond a certain demand level, clearly leading to a drop in speeds and a strong rise in the travel times. The space-time diagram in panel (b) shows the disruption of traffic flow caused by downward ToCs upstream of a No-AD zone in a single simulation for a specific parameter combination (LoD C, Mix 3).

To address the aforementioned impacts, [31] introduced several infrastructure-assisted traffic management measures designed for preventing, managing or distributing control transitions upstream of transition areas (areas on the roads where multiple control transitions may concurrently

1. CAVs will be informed about the existence of the upcoming No-AD zone via a simple information message, which causes automated vehicles to hand over control to the human driver at a specific position.

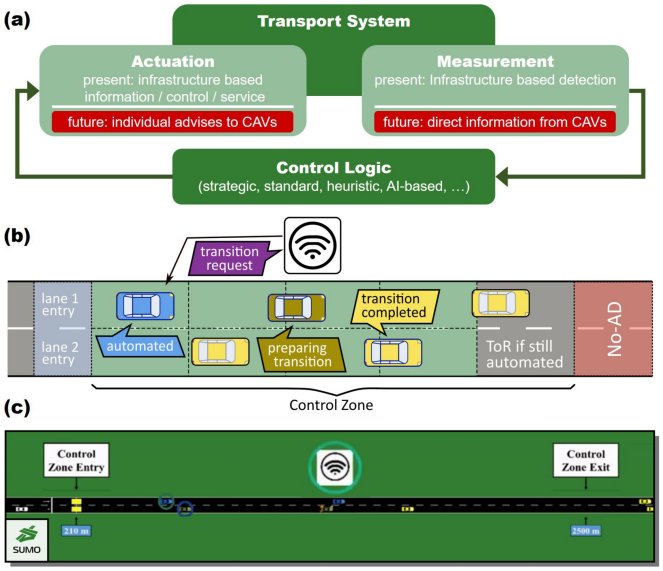

FIGURE 3. **Panel (a) shows the basic control scheme for a traffic management deploying CAVs for measurement and actuation in future traffic scenarios. Panel (b) illustrates a control scenario of a RSI communicating with automated vehicles within a defined control zone; colouring of vehicles corresponds to driving status; black dashed lines indicate the control cell borders. Panel (c) shows a SUMO screenshot of the discussed scenario.**

take place). Simulation findings showed that the proposed measures could mitigate the adverse impacts of control transitions for specific fleet mix and traffic demand scenarios. In particular, the distribution of downward ToCs in space and time upstream of a No-AD zone based on the TransAID approach could in many cases prevent traffic disruptions. The proposed distributed scheduling of downward ToCs reduced the local accumulation of accelerations/decelerations which may lead to strong speed variations and can consequently generate unsafe conditions. Thus, the rise in travel times observed in Fig. 2 could be at least postponed to a higher demand, or, in the best case, be avoided at all by AV specific ToR scheduling. Subsequently, the management of downward ToCs and guidance of MRMs to safe spots has been explored via real-world testing by Coll-Perales *et al.* in [32]. Their results suggested that the provision of personalized advice to CAVs including information about recommended ToC and safe spot locations could minimize MRMs taking place in lane which can result in hazardous situations.

B. REINFORCEMENT LEARNING IN TRAFFIC CONTROL

RL is a branch in machine learning that is heavily researched in emerging trends of deep learning applications nowadays, partly inspired by recent influential publications from [33] and [34]. Intelligent transportation systems (ITS) with often highly complex control tasks and vast volumes of data are a particularly promising field for data-driven techniques to improve nonlinear models as well as developing novel ideas to tackle future transportation challenges. The deployment of RL in traffic research was propelled by its ready utilization on control tasks in connection with traffic simulators. There is a wide range of traffic simulations that emulate a

certain system behavior. These representations of dynamic traffic systems conveniently serve as the environment component in RL. The feasibility to train an agent that operates and handles specific tasks in such environments allowed the rapid testing and progress of RL related techniques in traffic control. As [35] point out in their survey, various traditional transportation problems like demand or destination prediction, travel time estimation, traffic signal control or traffic flow prediction are investigated with help of deep learning and RL techniques.

Other areas of research and application where deep learning and RL have proven useful are traffic signal control [36], [37], [38] connected automated vehicles in mixed autonomy traffic [39], variable speed limit control at bottlenecks and ramps [40], or at round-abouts [41], to name but a few.

III. EXPERIMENTS

A. COMPUTER EXPERIMENTAL SETUP

We define the main control task as follows: A Traffic Management Center has to address vehicles within a confined area to prevent them from entering a No-AD zone,[2] still driving in automated mode. Therefore, a generic two-lane road is divided into two regions: (1) an upstream area denominated the control zone where automated driving is allowed and (2) a downstream area, where manual driving is mandatory (No-AD zone). Thus, the TMC issues a ToR to every CV and CAV approaching the No-AD zone. Vehicles that receive a ToR, initiate a downward ToC. A downward ToC results, for a certain time-span, in a moderate vehicle deceleration caused by the gap enlargement in the ToC preparation phase, and possibly reduced human driver performance. The control zone is managed by the TMC via V2X communication, e.g., using roadside infrastructure (RSI). The RSI receives position, speed and driving status from each CV and CAV via cooperative awareness messages [42] and maneuver coordination messages [43].

Fig. 3, panel (a) illustrates the envisioned role of connected vehicles within the standard control scheme of traffic management, actively adding sensory information and direct actuation to the control loop. Panel (b) details this simple idea of an RSI communicating with vehicles within the control zone by sending individual ToRs, aiming to prevent those vehicles to exit the control zone while still driving in automated mode. Panel (c) shows a snapshot of such a simulation experiment in SUMO.

For simplicity, failing downward ToCs resulting in MRMs are not managed specifically in this scenario.[3] Although MRMs are included in the simulations with a very rare average rate of occurrence, we do assume that human drivers

2. In this study we do not further examine the rationale for the mere existence of a No-AD zone. We presume the No-AD zone as a prerequisite for the TMC to address automated vehicles not to enter this area while driving automatically.

3. We refer the interested reader to [44] for a concept of how MRMs could be guided to safe spots if the given infrastructure permits.

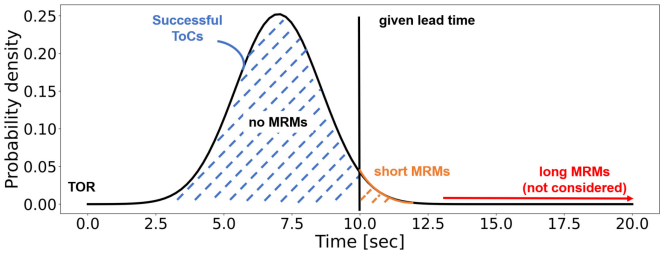

FIGURE 4. Reaction time probability distribution of the scenario.

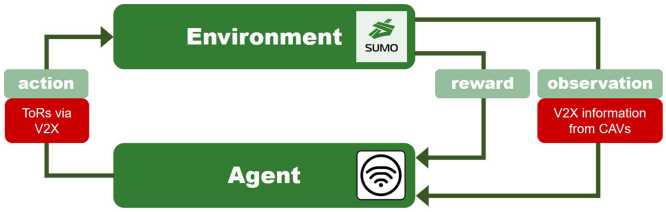

FIGURE 5. Interaction between agent and environment in a MDP (green boxes) including the prospective role of connected automated vehicles in this context (red boxes).

resume the vehicle's operation timely and do not block one lane, or both lanes, for an extended period of time. In the majority of the events, the driver takes over before the vehicle has stopped completely. Thus, occurring MRMs still have a negative impact on the traffic flow, but from a macroscopic perspective, they play only a secondary role relative to the by far more numerous downward ToCs. Fig. 4 shows the respective probability distribution of the reaction times by human drivers to take back control of the vehicle after receiving a ToR.

B. FORMULATION OF MARKOV DECISION PROCESS (MDP)

An MDP is a formalization of the decision making process of an agent that interacts with its environment. Specific actions by the agent affect the state of the environment. Based on observing the current state of the environment, the agent can take actions that aim to maximize future rewards, where rewards represent desirable outcomes. Fig. 5 schematically shows such an interaction.

In reference to Fig. 3, the environment represents the transport system, in our case emulated with the traffic simulator SUMO, and the agent corresponds to the traffic control logic. The data obtained by measurement and the means of interaction via V2X technology provide a basis for defining specific actions and observations in the MDP in this context. RL uses this formalization in order to guide and improve the decision making of an agent in such an interactive environment.

In this work we train a policy with Twin-Delayed Deep Deterministic Policy Gradient (TD3) [45], using the implementation of the *stable-baselines* library (version 2.10) [46]. A TD3 algorithm, which is a modification of Deep Deterministic Policy Gradient (DDPG) [47], simultaneously learns a Q-function for value updates and a

target-policy that maximizes the Q-function through gradient ascent. The interplay between those two is referred to as an actor-critic algorithm. TD3 improves on DDPG, by adding three tricks: (1) estimating the target with two Q-functions (clipped double Q-Learning), (2) employing a lower update rate on the target network (delayed policy update) and (3) adding random noise to the target policy (target policy smoothing) [47].

In the following, we formulate two models, *model 1* and *model 2*, which differ in observation space and reward function. The definitions for observations, actions and rewards in the next paragraphs are derived from initial work in [48].

1) OBSERVATIONS

The standard RL approach requires that the space of all possible states is of a fixed dimension. Therefore, the description of a dynamic vehicle flow as a list of vehicle states is not suitable, as this list greatly varies in length over time. We propose to overcome this problem by dividing the control zone of the highway into a constant number of cells per lane. In this manner, the agent can operate on a constant state space, despite facing a dynamic number of vehicles. The cells represent smaller parts of each lane with approximately the same size (cf. Fig. 3, black dashed lines indicate cell borders). We found that a number of $N_{\text{cells}} = 14$ cells represents a good trade-off between complexity and granularity for the scenario at hand. Every cell is described by three respectively four values, which constitute the perception of the environment for the agent:

(i) the average speed of all vehicles in the cell,
(ii) the number of manually driven vehicles (MVs) in the cell,
(iii) the number of CAVs in the cell,

For *model 1*, these quantities constitute the complete state of the environment. For *model 2*, we further provide

(iv) the number vehicles that are about to enter the control zone in the next time steps (see paragraph *c*) below for details).

2) ACTIONS

Similarly, each cell represents a potential target object for an action of the agent, i.e., the TMC. Effectively, an action would correspond to the transmission of ToRs to all CAVs in a specific cell at a specific time step. However, formally we construct the action in the MDP framework as a vector

$$\boldsymbol{a} \in [0, 1]^{N_{\text{cells}}},$$

which assigns a probability a_i, with $0 \leq a_i \leq 1$, to each cell, with which an effective action is triggered. Thus, at each control step (that is every second), the agent updates the probability for a ToR to every cell, taking into account the observed state. In this context, an assignment $a_i = 1$ corresponds to the deterministic decision of sending ToRs to cell i, while $a_i = 0$ ensures that no ToR is emitted.

Furthermore, the protocol ensures that a vehicle, which passes one of the last cells before the No-AD zone, will always receive a ToR, if it hasn't already. This is to avoid the vehicle from entering the No-AD zone while still driving automatically.

3) REWARD

At each control step, a reward

$$r(t) = r_{\text{ToR}}(t) + r_v(t), \tag{1}$$

is calculated based on a part $r_v(t)$ associated to the state of the environment and a part $r_{\text{ToR}}(t)$ associated to the number of effective ToRs transmitted by the TMC. For both considered models, we employ a ToR reward

$$\begin{aligned} r_{\text{ToR}}(t) = &\sum_{i=1}^{N_{\text{cells}}} r_{\text{ToR},i} \cdot \#\{\text{ToRs sent to cell } i\} \\ &- \pi_{\text{ToR}} \cdot \#\{\text{ToRs sent to last cells}\}. \end{aligned} \tag{2}$$

Here, $r_{\text{ToR},i}$ is a constant which defines the reward associated to a transmission of a ToR to a CAV in the i-th cell, which increases linearly from zero to w_{ToR} with the cell's index along its lane. That means, the further downstream a CAV is located when receiving its ToR, the higher the associated reward. As an exception, if the ToR is sent just before the No-AD zone, a penalty π_{ToR} is imposed, since a belated transmission increases the risk for CAVs to enter the No-AD zone in automated mode.

The fraction $r_v(t)$ of the reward (1) applied in *model 1* is proportional to the average speed $\bar{v}(t)$ taken over all vehicles in the control zone:

$$r_{v,\text{model1}}(t) = \bar{v}(t)/v_{\max}, \tag{3}$$

where $v_{\max} = 36m/s$ is the maximal allowed speed in the scenario.

In *model 2* this definition is extended by terms accounting for vehicles loaded into the simulation, but not entered, yet. For clarity, given a specific demand level, the simulation software SUMO generates vehicles at a corresponding rate and tries to insert these into the simulation scenario. If there is not sufficient free space on the road, it keeps the generated vehicles in a buffer. That means, these "pending" vehicles represent a tailback not depicted in the simulation and contain information regarding the traffic situation, valuable to its evaluation. Accordingly, we define the reward for *model 2* as

$$r_{v,\text{model2}}(t) = \bar{v}(t) + \bar{v}_{\text{pend}}(t) - \pi_{\text{pend}} \cdot \#\{\text{pending vehicles}\}, \tag{4}$$

where $\bar{v}_{\text{pend}}(t)$ is the average speed of all vehicles in the control zone *and* in the simulation buffer (accounted for with speed zero), and π_{pend} is a constant scaling the penalty imposed per loaded vehicle, not yet inserted in the simulation.

TABLE 1. Demand levels.

	Level of demand (LoD)		
	A	***B***	***C***
Q_{in} [veh/h]	1470	2310	3234

TABLE 2. Traffic compositions with vehicles shares for three different mixes.

Traffic mix	Vehicle Type		
	MV	***CV***	***CAV***
Mix 1	70%	15%	15%
Mix 2	50%	25%	25%
Mix 3	20%	40%	40%

TABLE 3. Vehicle types in the simulation represented by SUMO model combinations. The Krauß model is SUMO's standard model, while the ACC model is described in [49].

Driving Mode	SUMO Model	Vehicle Type		
		MV	***CV***	***CAV***
Car Following	Krauß	o	-	-
	ACC	-	o	o
Lane Change	Default	o	-	-
	Parametrized LC	-	o	o
Control Transition	ToC	-	o	o

C. SIMULATION AND TRAINING SETUP

We conducted our simulations with the microscopic traffic simulator SUMO [25], version *1.6*. On a two-lane motorway with a length of 5.0 km and a speed limit of 130 km/h, vehicles enter the network randomly with a Poissonian distribution with a demand Q_{in} at the upstream part of the road. The maximum capacity of the two-lane motorway is assumed to be 4200 $[veh/h]$ for a homogeneous fleet of MVs, that means without the existence of control transitions. Since the maximum road capacity inevitably decreases in the presence of control transitions in all cases, and also in every case differently, we compare those TM approaches based on the induced demand. Three different demand levels *LoD* were defined (see Table 1).

The following tables summarize the three different traffic mixes with shares of manually driven (MV), connected (CV), and connected automated vehicles (CAV) used in the experiment (see Table 2), and also the respective vehicle models represented by the different models in SUMO (see Table 3). The selection of the simulated traffic mixes was made to verify that increasing shares of CV/CAVs escalate traffic disruption due to higher frequency of downward ToC events which are not efficiently distributed in space and time, rather than to explicitly quantify the exact penetration rates of CVs/CAVs inducing traffic flow breakdown as a result of accumulated downward ToCs. In the context of our simulation experiments, we consider that drivers of CVs continuously monitor the operation of the automation functions and can promptly react to ToRs. Moreover, CVs are ACC/CACC capable and their lane change behavior is more conservative compared to MVs. On the other hand, drivers of CAVs can be involved in secondary tasks during normal operation in automated mode, and thus exhibit delayed response to ToRs or even fail to resume vehicle control within the available lead time (CAV executes MRM on these rare occasions). CAVs are also assumed ACC/CACC capable, but their lane change behavior is more conservative compared to CVs. A detailed parametrization of the utilized SUMO models per vehicle type can be found in [30].

In the simulation, the No-AD zone starts at 2.5 km downstream of the network entry. The TM controller addresses CVs and CAVs ahead of the No-AD zone by sending take-over requests via the SUMO API *traci*.

For the parametrization of the TD3 training we mainly use the default parameters provided by the *stable baselines* library. For approximating the policy and the Q-function we deploy neural networks with two hidden layers, using [300, 400] for layer size as in [47]. The only hyperparameters, which have been altered in favor of the best results, are:

1) *buffer_size*: The size of the replay buffer to save experiences at each step, to update the target network episodically was set to 100000.
2) *train_freq*: The value to define the model update rate of the target network was set to 300 steps.

Two different models termed *model 1* and *model 2* were trained with the respective rewards functions. For the respective reward parameters we chose: $w_{\text{ToR}} = 10$, $\pi_{\text{ToR}} = 100$, $\pi_{\text{pend}} = 1/1200$. A training episode ran for 1200s, with the first 200s as uncontrolled warm-up in order to establish a fully populated control zone. We let the model learn for 2000 episodes with a random seed for each episode. Both policies were trained on the highest demand level *C* with the vehicle fleet of *Mix 3*.

After the finished training, we ran simulations with both models for all parameter combinations (LoD/Mix), each combination with 10 random runs for 1 hour simulated time per seed.

IV. RESULTS

In the following we present the results obtained from the simulation study. We ran simulations for all parameter combinations (LoD/Mix) for *model 1*, *model 2*, the do-nothing case (*baseline*), the deterministic control approach proposed by the *TransAID* consortium [31], and a *random* model, that encompasses uniformly distributed ToRs within the control zone. Since our RL-based models were explicitly designed with the consideration of mobility objectives, we focus on the analysis of simulation results from the traffic efficiency perspective. First, we evaluate the performance of all cases based on the average travel times per vehicle and the average distance driven in automated mode (CAV distance). Also, we compare the control strategies for both trained models based on space time diagrams and the spatial distribution of sent ToRs. Finally, we analyze the training success by evaluating the reward.

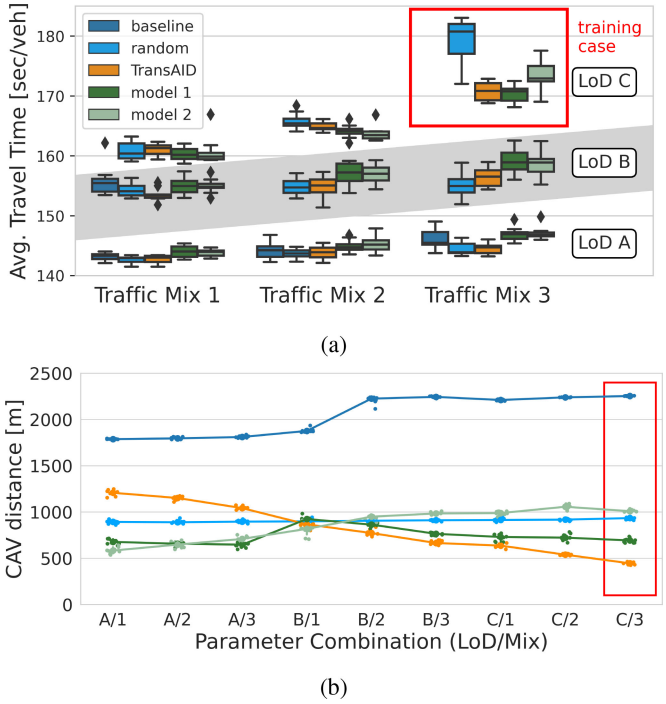

FIGURE 6. **Aggregated results for five different cases/models: *baseline* (dark blue), *random* (light blue) *TransAID* (orange), *model 1* (dark green), *model 2* (light green). Panel (a): Travel time for all parameter combinations; results of the baseline simulation for parameter combinations higher than B1/Mix 1 are excluded. The medians for those excluded combinations in ascending order are: 620.3, 978.2, 457.6, 721.8 and 1093.6*s*. Compare also to Fig. 2, panel (a). Panel (b): Average CAV distance driven in automated mode. The lines are drawn to guide the readers eyes.**

A. MOBILITY PERFORMANCE

Fig. 6 presents the aggregated results of the simulation study for all parameter combinations for both models compared to the unmanaged (baseline) and also two alternatively managed cases (TransAID) and (random). Panel (a) shows the individual vehicle's average travel times in the form of boxplots. Panel (b) respectively presents the average covered distance of CAVs within the control zone driving in automated mode.[4]

Firstly, comparing only the two trained models with each other, we observe that travel times are almost the same for all combinations except for C/3 where *model 1* slightly outperforms *model 2*. Notably, this is the combination the models were trained with (cf. red boxes in Fig. 6). On the other hand, for C/3 *model 2* covers significantly more CAV distance than *model 1* (about plus 300m per vehicle on average, which corresponds to approximately one cell length), corresponding to the rationale of the model development, see Section III-B. For parameter combinations C/2 to B/1 *model 2* is able to prolong automated driving for CAVs longer than *model 1*. For LoD A differences in CAV distance are rather small. So, the plain side-by-side comparison of both models leads to the initial conclusion that both achieve a good compromise between optimizing travel time and CAV distances simultaneously, yet with different emphasises on favouring travel time vs. CAV distance (cf. Fig. 6, case C/3).

The good performance of both models becomes obvious at higher demand levels when compared to the the other three approaches. For the baseline, we first take a look at panel (b) with the CAV distance. Due to the fact that all CAVs perform the control transition shortly before entering the No-AD zone, they cover the maximum possible CAV distance, but since consecutive and simultaneous downward ToCs cause disruptions in traffic flow, the travel time in panel (a) increases significantly (up to the factor 6, see Fig. 2) which ultimately results in traffic jams for parameter combinations higher than B/1.[5] In contrast, the original approach from TransAID is based on distributing individual ToRs, by forming virtual platoons for consecutive CAVs. This allows a continuous optimization of the latest possible ToR depending on the current speed of a CAV, achieving quite low average travel times, in fact slightly better than *model 2* and similar compared to *model 1* (see panel (a)). But, since this approach strongly depends on the traffic density within the control zone, in panel (b) we can observe a steady decrease in CAVs distance for higher parameter combinations. Both models prevent this decline for LoD B and LoD C, but drop noticeably in CAV distance for LoD A.

Comparing the random case to both models, it is noticeable, that this rather simple heuristic approach performs very well for LoD A and LoD B in terms of travel time and CAV distance, similar to *TransAID*. For C/3 though, this random approach apparently hits the capacity limit way earlier and shows significantly higher travel times.

Overall, the highest benefits of the two RL models compared to all other cases can be gained at LoD C depending on the performance metric (travel time vs. CAV distance). Both models seem to be able to handle lower demands and mixes, i.e., LoD B/Mix 2 and Mix 3, that are somewhat close to the training case, but for LoD A that mobility performance drops for the RL models.

B. CONTROL STRATEGY PERFORMANCE

In Fig. 7 two exemplary space time diagrams are shown, which correspond to the same parameter combination as in Fig. 2. They illustrate that both models are very well able to distribute downward ToCs so that no traffic jams develop. The visible differences between panel (a) and (b) in the distribution of the mean speed within the control zone (the area below the white dashed line) result from the different control strategies of the two models. The areas of lighter blue indicate the occurrence of short episodes of slower

4. Note, for better visibility, in panel (a) we excluded boxes of the baseline simulation for parameter combinations higher than B1/Mix 1. In panel (b) we used grouped scatter plots and drew additional lines between groups for each model.

5. Notice that for the baseline, the average covered CAV distance for demand level LoD A and B1/Mix1 is less than 2000m opposed to the rest of the combinations with 2400m. This is because the ToR message is triggered dependent on the current vehicle speed. For combinations with low travel times, vehicles can drive almost up to their desired speed, which is about $36m/s$ opposed to the rest of the parameter combinations, where vehicles in congested traffic only drive about $5 - 10m/s$.

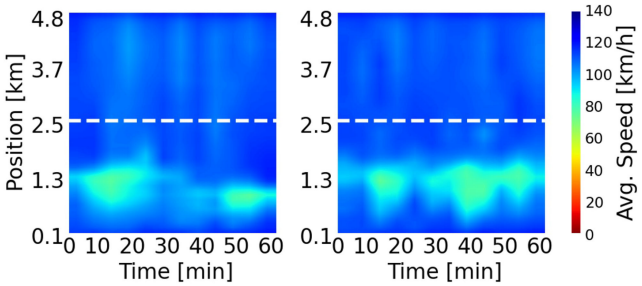

(a) *model 1* (b) *model 2*

FIGURE 7. **Spatiotemporal diagrams of mean speed for a simulation run with *model 1* (left) and *model 2* (right) of a single simulation run: LoD C - Mix 3 - Seed 7. The white dashed line represents the point from which all vehicles have to drive manually.**

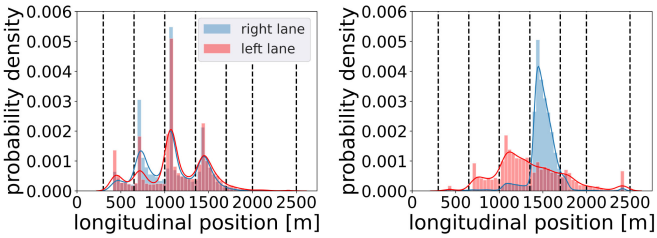

(a) model 1 (b) model 2

FIGURE 8. **Spatial probability distribution of take-over requests sent to CV/CAVs within the range of the control zone for *model 1* (left) and *model 2* (right) for parameter combination LoD C/Mix 3. The black dashed vertical lines mark the cell edges defined in Section III-B, showing 7 cells per lane.**

average speed, which do not build up to a persistent congestion, though. Also, they support the argument that the TMCs objective is to preserve a rather smooth traffic flow without too many disruptive decelerations/accelerations by consecutive vehicles at similar positions. Clearly, *model 1* and *model 2* improve the smoothness of the traffic flow in comparison with the unmanaged case shown in Fig. 2. Thus, the rationale of the reward design seems to have identified mechanisms related to the emergence of congestion in that case. The trained controllers manage successfully and counteract congestion by a distribution of ToRs.

However, the spatiotemporal diagrams rather indirectly point to the control strategy, since the vehicle speed is a delayed indicator of the TMCs distribution strategy. Therefore Fig. 8 shows the spatial distribution of ToRs sent to CAVs within the control zone for each lane. As a clear distinction, we observe a disparate utilization of the left and right lane between both models, as well as different spatial distributions over the length of the control zone. Whereas *model 1* uses both lanes almost equally close to a normal distribution pattern around 1050*m*, *model 2* shows significant differences in distribution (skewed distribution on left lane) and lane utilization shifted further downstream (1600*m* on right lane). This might be connected to differences in the overall performance of the control strategies. On the contrary, both models show similarities in not sending ToRs to the last cells before the No-AD zone (longitudinal position >2000m), as imposed by the reward definitions.

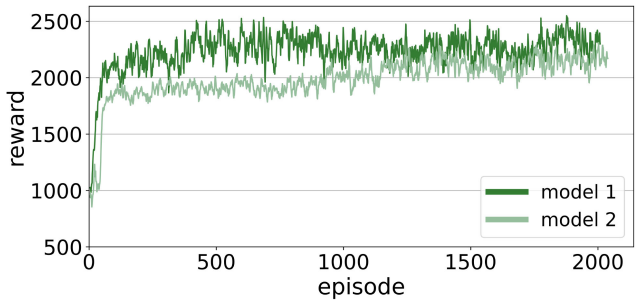

FIGURE 9. **Reward per episode for *model 1* and *model 2*.**

Despite these differences, the ToC management performs well in both cases while achieving a compromise between travel time optimization versus preserving the automated driving mode.

C. TRAINING ASSESSMENT

Fig. 9 shows the reward curves for the training process of the RL-based models 1 and 2. Both models converge relatively quickly, although *model 2* apparently needs more episodes to do so. Given that we effectively trained for 1000 steps per episode, the final average return is about 2.2 - 2.4 for both models. However, note that the average return for *model 2* may be expected to deviate slightly due to the different definition for $r_v(t)$ in Eq. (1), cf. (4).

Overall, some of the variance in the reward could be explained due to the randomness in vehicle flow and numbers. For the training of models 1 and 2, leading to the performance reported in Section IV, we ran each training run for 2000 episodes with random seeds per episode.

Additionally, we tested training runs with fixed seeds, i.e., deterministic vehicle flows for each episode, in order to find a smoother reward convergence (as e.g., in [50]). Those test runs (not displayed here) indeed showed less variance in the reward and also converged at about the same level as the training with random seeds. Nonetheless, the random seed training induced a better overall performance. We consider this to be attributable to an increased versatility of the control gained from the confrontation with a larger variety of situations and allowing it to handle different LoD and traffic mixes in a more robust fashion.

Moreover, it is worthwhile mentioning that the choice of the reward definition is critical for the training success, as well. For example, during the development, we tested one function, which imposed rather strong penalties (i.e., negative rewards) on the sending of ToRs. The penalties included a similar cost gradient as in (2), which favored longer distances in automated driving mode. From this definition, the model did not learn to postpone ToR transmissions, as intended. In contrast, it rather learned to reduce the occurrence of ToRs via a reduction of the inflow. This was achieved by sending all ToRs to the first cell in contradiction to the original objective of prolonging the distance of automated driving.

Clearly, an accurate design of the reward is necessary to align the resulting model behavior with the training objectives.

V. CONCLUSION

For a range of different demand levels and ratios of automated and manually driven vehicles, we have demonstrated, that traffic control can significantly attenuate the negative side-effect of control transitions from automated to manual. Four different TM methods have been explored and compared to a do-nothing scenario: Two heuristic approaches adhering to the form of conventional TM protocols, and two based on a RL approach employing slightly different reward functions. All methods perform similarly well, in terms of travel times and slight improvements with the trained RL models for prolonging the automated driving mode. Furthermore, all protocols outperform a do-nothing solution. Especially for higher demand levels the performance gap is significant. It is noteworthy that the RL models were trained with just one, relatively high demand at a fixed, relatively high ratio of automated vehicles in the traffic composition, but still perform properly when getting applied to lower demands and other fleet mixes in our scenario.

The robustness of the RL approaches has been tested by changing the demand and the vehicle fleet composition. In all cases, we see that the traffic management can ameliorate the potential capacity drop induced by an accumulation of downward ToCs in transition areas. We assess that this delay of the foreseeable capacity drop under high demand and with high CAVs percentages will be the task at hand for future TMCs in comparable scenarios. However, for very large demands close to capacity limit all of the methods must finally fail. While the two RL models, as well as the conventional TM controller, do not differ considerably in the maximum capacity they finally achieve, they enhance the capacity significantly compared to the baseline approach. The conventional TM controller *TransAID*, although way better than a do-nothing approach, manages high demand levels only by compromising the objective of preserving the automating driving mode. On the contrary, the naive random distribution TM approach achieves better CAV distances than *TransAID*, but only performs adequately for lower demands and hits the capacity limit earlier than the RL models.

We acknowledge that the sole objective to maintain automated driving as long as possible, although a downward ToC cannot be avoided in such a scenario, might be conflicting from a traffic management perspective and should only be considered while simultaneously tackling the adverse impacts discovered in the baseline analysis. Manufacturers and consumers of CAVs on the one hand might be interested in maintaining automated driving features without external interference, while from a traffic safety and efficiency perspective it can be favourable to initiate downward ToCs further upstream. In denser and highly heterogeneous traffic, belated ToRs resulting in ToCs close to a No-AD zone also mean more complex and possibly unsafe interactions between MVs and CV/CAVs. Therefore preserving a smooth and safe traffic flow should be a priority in such scenarios for a TMC. Although limitations of the ToC model, such as not capturing evasive and overtaking manoeuvres, may amplify the adverse downward ToC impacts discussed before, we think that the results presented illustrate the feasibility to accomplish both objectives (efficient traffic + automated driving) concurrently as long as demand and traffic mix do not exceed the capacity limit in the scenario. In that regard, the RL method achieves a better overall performance.

Accordingly, this work shows that an AI-based approach is very well able to be on par with more traditional approaches, which usually require a lot of traffic domain knowledge to be developed. Thus, AI-approaches open the promising perspective of control solutions with a limited need for such expertise, and offering a potential for synergies with developments in other areas. It is still a challenging and time-consuming task, though, to develop adequate RL models, which usually require a careful tailoring to the given problem. Besides, the duration of the development cycle of adapting the MDP setup and assessing the learning progress of the altered model is often governed by the computational complexity of the task. In our case, the training of a model variant took about 24 hours (on a Intel Core i9-10900X CPU 10×3.7GHz) and numerous iterations were necessary to find an adequate formalization and parametrization of the training experiment.

Finally, we plan to adapt our RL-based models in future work so that they also account for safety objectives, and conduct an explicit microscopic traffic simulation based safety evaluation that will encompass an in-depth analysis of conflicts and relevant surrogate safety assessment measures (SSMs). Moreover, the aspect on how to handle MRMs in such traffic scenarios, which were mostly ignored in this study, might be a future research focus for applying RL to elaborate TM schemes, but it will take significant effort and time, considering the rather low sample efficiency of the training process. In addition, we think that not much more can be gained in terms of traffic efficiency with any other TM approach, for a scenario such as we presented, since a control transition itself, as modelled here, diminishes capacity, and there is no method that can bring it back.

ACKNOWLEDGEMENT

Open access publication fees were covered by the DLR Publication Fund.

REFERENCES

[1] On-Road Automated Driving (ORAD) Committee (SAE Int., Warrendale, PA, USA). *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*, (Jun. 2018). [Online]. Available: https://doi.org/10.4271/J3016_201806

[2] F. Favaró, S. Eurich, and N. Nader, "Autonomous vehicles disengagements: Trends, triggers, and regulatory limitations," *Accid. Anal. Prevent.*, vol. 110, pp. 136–148, Jan. 2018. [Online]. Available: https://doi.org/10.1016/j.aap.2017.11.001

[3] Z. Lu, R. Happee, C. D. D. Cabrall, M. Kyriakidis, and J. C. F. de Winter, "Human factors of transitions in automated driving: A general framework and literature survey," *Transp. Res. F, Traffic Psychol. Behav.*, vol. 43, pp. 183–196, Nov. 2016. [Online]. Available: https://doi.org/10.1016/j.trf.2016.10.007

[4] V. Melcher, S. Rauh, F. Diederichs, H. Widlroither, and W. Bauer, "Take-over requests for automated driving," *Procedia Manuf.*, vol. 3, pp. 2867–2873, Jan. 2015. [Online]. Available: https://doi.org/10.1016/j.promfg.2015.07.788

[5] S. Petermeijer, P. Bazilinskyy, K. Bengler, and J. de Winter, "Take-over again: Investigating multimodal and directional tors to get the driver back into the loop," *Appl. Ergonom.*, vol. 62, pp. 204–215, Jul. 2017. [Online]. Available: https://doi.org/10.1016/j.apergo.2017.02.023

[6] M. Bahram, M. Aeberhard, and D. Wollherr, "Please take over! an analysis and strategy for a driver take over request during autonomous driving," in *Proc. IEEE Intell. Veh. Symp. (IV)*, 2015, pp. 913–919, doi: 10.1109/IVS.2015.7225801.

[7] S. A. M. Agriesti, M. Ponti, G. Marchionni, and P. Gandini, "Cooperative messages to enhance the performance of L3 vehicles approaching roadworks," *Eur. Transp. Res. Rev.*, vol. 13, p. 1, Jan. 2021. [Online]. Available: https://doi.org/10.1186/s12544-020-00457-z

[8] R. Bellman, "A Markovian decision process," *J. Math. Mech.*, vol. 6, no. 5, pp. 679–684, 1957. [Online]. Available: http://www.jstor.org/stable/24900506

[9] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018. [Online]. Available: http://incompleteideas.net/book/the-book-2nd.html

[10] M. Coşkun, A. Baggag, and S. Chawla, "Deep reinforcement learning for traffic light optimization," in *Proc. IEEE Int. Conf. Data Min. Workshops (ICDMW)*, Singapore, 2018, pp. 564–571, doi: 10.1109/ICDMW.2018.00088.

[11] F. Belletti, D. Haziza, G. Gomes, and A. M. Bayen, "Expert level control of ramp metering based on multi-task deep reinforcement learning," *IEEE Trans. Intell. Transp. Syst.*, vol. 19, no. 4, pp. 1198–1207, Apr. 2018, doi: 10.1109/TITS.2017.2725912.

[12] Q. Guo, O. Angah, Z. Liu, and X. J. Ban, "Hybrid deep reinforcement learning based eco-driving for low-level connected and automated vehicles along signalized corridors," *Transp. Res. C, Emerg. Technol.*, vol. 124, Mar. 2021, Art. no. 102980. [Online]. Available: https://doi.org/10.1016/j.trc.2021.102980

[13] D. M. Vlachogiannis, E. I. Vlahogianni, and J. Golias, "A reinforcement learning model for personalized driving policies identification," *Int. J. Transp. Sci. Technol.*, vol. 9, no. 4, pp. 299–308, 2020. [Online]. Available: https://doi.org/10.1016/j.ijtst.2020.03.002

[14] D. Nallaperuma *et al.*, "Online incremental machine learning platform for big data-driven smart traffic management," *IEEE Trans. Intell. Transp. Syst.*, vol. 20, no. 12, pp. 4679–4690, Dec. 2019, doi: 10.1109/TITS.2019.2924883.

[15] "Easy Tech: Conditionally Automated Driving With the Drive Pilot." Jul. 2021. [Online]. Available: https://www.daimler.com/magazine/technology-innovation/easy-tech-drive-pilot.html (Accessed: Dec. 13, 2021).

[16] F. M. Favaró, N. Nader, S. O. Eurich, M. Tripp, and N. Varadaraju, "Examining accident reports involving autonomous vehicles in California," *PLoS One*, vol. 12, pp. 1–20, Sep. 2017. [Online]. Available: https://doi.org/10.1371/journal.pone.0184952

[17] S. Wang and Z. Li, "Exploring causes and effects of automated vehicle disengagement using statistical modeling and classification tree based on field test data," *Accid. Anal. Prevent.*, vol. 129, pp. 44–54, Aug. 2019. [Online]. Available: https://doi.org/10.1016/j.aap.2019.04.015

[18] A. M. Boggs, R. Arvin, and A. J. Khattak, "Exploring the who, what, when, where, and why of automated vehicle disengagements," *Accid. Anal. Prevent.*, vol. 136, Mar. 2020, Art. no. 105406. [Online]. Available: https://doi.org/10.1016/j.aap.2019.105406

[19] Z. H. Khattak, M. D. Fontaine, and B. L. Smith, "Exploratory investigation of disengagements and crashes in autonomous vehicles under mixed traffic: An endogenous switching regime framework," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 12, pp. 7485–7495, Dec. 2021. [Online]. Available: https://doi.org/10.1109/TITS.2020.3003527

[20] A. D. McDonald *et al.*, "Toward computational simulations of behavior during automated driving takeovers: A review of the empirical and modeling literatures," *Human Factors J. Human Factors Ergonom. Soc.*, vol. 61, no. 4, pp. 642–688, 2019. [Online]. Available: https://doi.org/10.1177/0018720819829572

[21] X. Xin *et al.*, "A literature review of the research on take-over situation in autonomous driving," in *Proc. Int. Conf. Human-Comput. Interact.*, 2019, pp. 160–169, doi: 10.1007/978-3-030-23538-3_12.

[22] J. W. Van Lint and S. C. Calvert, "A generic multi-level framework for microscopic traffic simulation—Theory and an example case in modelling driver distraction," *Transp. Res. B, Methodol.*, vol. 117, pp. 63–86, Nov. 2018. [Online]. Available: https://doi.org/10.1016/j.trb.2018.08.009

[23] S. C. Calvert and B. van Arem, "A generic multi-level framework for microscopic traffic simulation with automated vehicles in mixed traffic," *Transp. Res. C, Emerg. Technol.*, vol. 110, pp. 291–311, Jan. 2020. [Online]. Available: https://doi.org/10.1016/j.trc.2019.11.019

[24] L. Lücken, E. Mintsis, K. Porfyri, R. Alms, Y.-P. Flötteröd, and D. Koutras, "From automated to manual—Modeling control transitions with sumo," in *Proc. SUMO User Conf.*, 2019, pp. 124–144. [Online]. Available: https://doi.org/10.29007/sfgk

[25] P. A. Lopez *et al.*, "Microscopic traffic simulation using SUMO," in *Proc. 21st IEEE Int. Conf. Intell. Transp. Syst.*, Nov. 2018, pp. 2575–2582. [Online]. Available: https://elib.dlr.de/127994/

[26] E. P. Todosiev, "The action point model of the driver-vehicle system," Ph.D. dissertation, Dept. Doctor Philos., Ohio State Univ., Athens OH, USA, 1963. [Online]. Available: http://rave.ohiolink.edu/etdc/view?acc_num=osu1486555089597102

[27] W. Xin, J. Hourdos, P. Michalopoulos, and G. Davis, "The less-than-perfect driver: A model of collision-inclusive car-following behavior," *Transp. Res. Rec.*, vol. 2088, no. 1, pp. 126–137, 2008. [Online]. Available: https://doi.org/10.3141/2088-14

[28] C. Gardiner, *Stochastic Methods: A Handbook for the Natural and Social Sciences*, 4th ed. Berlin, Germany: Springer, 2009. [Online]. Available: https://link.springer.com/book/9783540707127

[29] M. Treiber and A. Kesting, *Traffic Flow Dynamics: Data, Models and Simulation*, 1st ed. Heidelberg, Germany: Springer, 2013. [Online]. Available: https://link.springer.com/book/10.1007/978-3-642-32460-4#about

[30] E. Mintsis *et al.* "TransAID Deliverable 3.1—Modelling, Simulation and Assessment of Vehicle Automations and Automated Vehicles' Driver Behaviour in Mixed Traffic." 2019. [Online]. Available: https://cordis.europa.eu/project/id/723390/results

[31] S. Maerivoet *et al.* "TransAID Deliverable 4.2—Preliminary Simulation and Assessment of Enhanced Traffic Management Measures." 2019. [Online]. Available: https://cordis.europa.eu/project/id/723390/results

[32] B. Coll-Perales *et al.*, "Prototyping and evaluation of infrastructure-assisted transition of control for cooperative automated vehicles," *IEEE Trans. Intell. Transp. Syst.*, early access, Mar. 4, 2021. [Online]. Available: https://doi.org/10.1109/TITS.2021.3061085

[33] V. Mnih *et al.*, "Playing atari with deep reinforcement learning," 2013, *arXiv:1312.5602*.

[34] V. Mnih *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, pp. 529–533, Feb. 2015, doi: 10.1038/nature14236.

[35] M. Veres and M. Moussa, "Deep learning for intelligent transportation systems: A survey of emerging trends," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 8, pp. 3152–3168, Aug. 2020. [Online]. Available: https://doi.org/10.1109/TITS.2019.2929020

[36] T. Chu, J. Wang, L. Codecá, and Z. Li, "Multi-agent deep reinforcement learning for large-scale traffic signal control," *IEEE Trans. Intell. Transp. Syst.*, vol. 21, no. 3, pp. 1086–1095, Mar. 2020. [Online]. Available: https://doi.org/10.1109/TITS.2019.2901791

[37] J. V. S. Busch, V. Latzko, M. Reisslein, and F. H. P. Fitzek, "Optimised traffic light management through reinforcement learning: Traffic state agnostic agent vs. holistic agent with current V2I traffic state knowledge," *IEEE Open J. Intell. Transp. Syst.*, vol. 1, pp. 201–216, 2020. [Online]. Available: https://doi.org/10.1109/OJITS.2020.3027518

[38] H. Wang, H. Chen, Q. Wu, C. Ma, and Y. Li, "Multi-intersection traffic optimisation: A benchmark dataset and a strong baseline," *IEEE Open J. Intell. Transp. Syst.*, vol. 3, pp. 126–136, 2021, doi: 10.1109/OJITS.2021.3126126.

[39] E. Vinitsky *et al.*, "Benchmarks for reinforcement learning in mixed-autonomy traffic," in *Proc. 2nd Conf. Robot Learn.*, Oct. 2018, pp. 399–409. [Online]. Available: http://proceedings.mlr.press/v87/vinitsky18a.html

[40] Z. Li, P. Liu, C. Xu, H. Duan, and W. Wang, "Reinforcement learning-based variable speed limit control strategy to reduce traffic congestion at freeway recurrent bottlenecks," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 11, pp. 3204–3217, Nov. 2017. [Online]. Available: https://doi.org/10.1109/TITS.2017.2687620

[41] G. Bacchiani, D. Molinari, and M. Patander, "Microscopic traffic simulation by cooperative multi-agent deep reinforcement learning," 2019, *arXiv:1903.01365*.

[42] *Intelligent Transport Systems (ITS); Vehicular Communications; Basic Set of Applications; Part 2: Specification of Cooperative Awareness Basic Service*, ETSI Standard EN 302 637-2, Apr. 2019. [Online]. Available: https://www.etsi.org/deliver/etsi_en/302600_302699/30263702/01.04.01_60/en_30263702v010401p.pdf

[43] "Intelligent transport systems (ITS); vehicular communications; informative report for the maneuver coordination service, draft 0.0.4," ETSI, Sophia Antipolis, France, ETSI Rep. TR 103 578, 2019.

[44] R. Alms, Y.-P. Flötteröd, E. Mintsis, S. Maerivoet, and A. Correa, "Traffic management for connected and automated vehicles on urban corridors—Distributing take-over requests and assigning safe spots," in *Proc. 3rd Symp. Manag. Future Motorway Urban Traffic Syst.*, Jul. 2020, pp. 1–11. [Online]. Available: https://elib.dlr.de/134136/

[45] S. Fujimoto, H. van Hoof, and D. Meger, "Addressing function approximation error in actor-critic methods," 2018, *arXiv:1802.09477*.

[46] A. Hill *et al.* "Stable Baselines." 2018. [Online]. Available: https://github.com/hill-a/stable-baselines

[47] T. P. Lillicrap *et al.*, "Continuous control with deep reinforcement learning," in *Proc. 4th Int. Conf. Learn. Represent. (ICLR)*, San Juan, Puerto Rico, May 2016. [Online]. Available: http://arxiv.org/abs/1509.02971

[48] A. Noulis, "Reinforcement learning in traffic control for connected automated vehicles," M.S. thesis, Dept. Inst. Transp. Syst., TU Berlin, Berlin, Germany, Oct. 2020. [Online]. Available: https://elib.dlr.de/139169/

[49] V. Milanés and S. E. Shladover, "Modeling cooperative and autonomous adaptive cruise control dynamic responses using experimental data," *Transp. Res. C, Emerg. Technol.*, vol. 48, pp. 285–300, Nov. 2014. [Online]. Available: https://doi.org/10.1016/j.trc.2014.09.001

[50] E. Vinitsky, N. Lichtle, K. Parvate, and A. Bayen, "Optimizing mixed autonomy traffic flow with decentralized autonomous vehicles and multi-agent RL," 2020, *arXiv:2011.00120*.

3.3 Paper II - Control Transitions in Level 3 Automation: Safety Implications in Mixed-Autonomy Traffic

Originally published as:

Alms, R.; Wagner, P. "Control Transitions in Level 3 Automation: Safety Implications in Mixed-Autonomy Traffic," *Safety* **2024**, *10*(1), 1. DOI: 10.3390/safety10010001

Scope and Contribution of Paper II

Following the efficiency-focused analysis in Paper I, Paper II addresses RQ 2 by examining the safety implications of Level 3 ToCs in a different scenario. With ALKS approval limited to 60 km/h in the 2021 version of UNECE Regulation No. 157 [79], several automotive manufacturers introduced their first Level 3 systems to the market[6]. Motivated by this speed restriction in the ODD, a use case was defined to investigate potential implications for traffic safety[7]. Accordingly, Paper II develops a SUMO-based experiment contrasting a *Base* scenario (no ToCs) with a *Level 3* scenario (ToCs triggered at the ODD boundary) across multiple traffic mixes and demand levels up to and near capacity. Safety is assessed using surrogate safety measures (SSMs), namely time to collision (TTC) and modified deceleration rate to avoid a crash (MDRAC; a variant of DRAC, the deceleration rate to avoid a crash). The results indicate that consecutive ToCs in an accelerating vehicle stream at the ODD boundary (i.e. speeds ≥ 60 km/h) degrade safety near capacity once the Level 3 share becomes non-negligible, beginning with Mix 1 (20% AVs). Two additional experiment variants investigate potential mitigation and stability issues: an open-gap preparation setting, which suppresses active deceleration during ToC preparation, removes the detrimental pattern and suggests anticipatory headway management for AVs as an effective countermeasure; by contrast, a string-unstable ACC variant reveals amplified safety risks, indicating an adverse interaction between controller instability and ToC timing.

With respect to RQ 2, safety-relevant effects occur under near-capacity traffic

[6]Honda in a limited series in 2021 in Japan; Mercedes in 2023 in Germany and the US; BMW in 2024 in Germany.

[7]The amended version of Regulation No. 157 from 2023 [78], which extends the ODD range up to 130 km/h, and the subsequent approval of Mercedes' Level 3 system up to 95 km/h in 2024, will be discussed in Paper III.

when ToCs concentrate at the ODD boundary. Mitigation in such a speed-limited ODD (whilst under accelerating traffic conditions) is possible through an anticipatory headway element, either implemented via infrastructure-assisted traffic management (TM) or vehicle-system prediction. Section 4.2 augments these results with an additional ODD-limited scenario.

Correction note:

The following Figure 3.3 presents a corrected version of the use case illustration from Figure 3 in Paper II. Specifically, the original graphic displays a speed profile in panel (b) that does not correctly reflect the speed limits illustrated in panel (a), i.e., 60 km/h and 100 km/h. In Figure 3.3, the updated profile in panel (b) now accurately represents the increasing average velocity in the scenario after the speed limit increases from 60 to 100 km/h, illustrating the effects of mandatory ToCs in Mix 5 (including AVs) compared to Mix 0 (without AVs).[8]

We proposed this correction to the MDPI *Safety* journal, but the request was declined by the editorial office, as it was considered a minor issue.

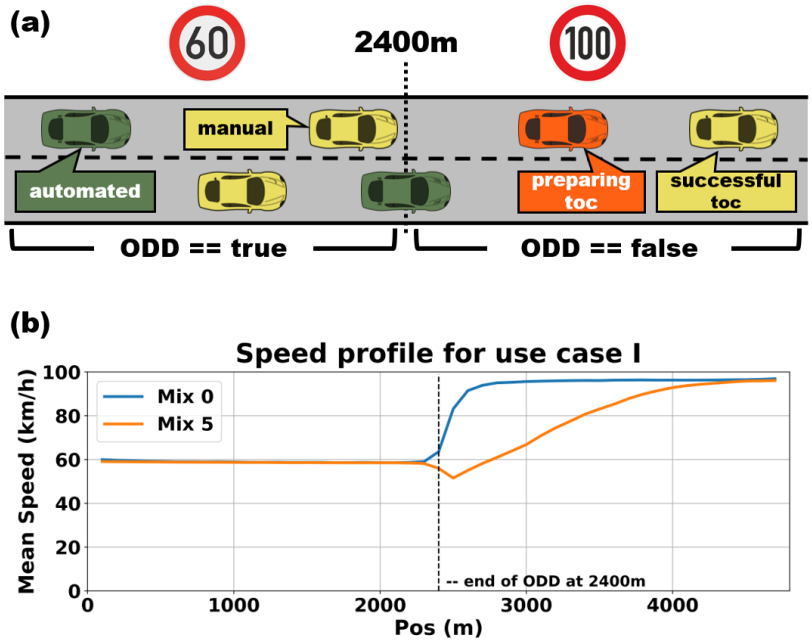

Figure 3.3: Corrected speed profile of the use case illustrated in Figure 3, panel (b).

[8]Note that due to a LGV/HGV share of 15% in these simulations, the average speed does not reach the full 100 km/h in the second section when the ODD is off. This is caused by the lower maximum speeds of such trucks, resulting in an average velocity slightly below the speed limit.

safety

MDPI

Article

Control Transitions in Level 3 Automation: Safety Implications in Mixed-Autonomy Traffic

Robert Alms [1,*] and Peter Wagner [1,2]

[1] Institute of Transportation Systems, German Aerospace Center (DLR), Rutherfordstr 2, 12489 Berlin, Germany
[2] Institute of Land and Sea Transport Systems, TU Berlin, Salzufer 17–19, 10587 Berlin, Germany
* Correspondence: robert.alms@dlr.de

Abstract: Level 3 automated driving systems could introduce challenges to traffic systems as they require a specific lead time in their procedures to ensure the safe return of vehicle control to the driver. These processes, called 'transitions of control', may particularly pose complications in accelerating traffic flows when regulations mandate control transitions due to an operational speed limitation of 60 km/h as established in recent certification processes based on UNECE regulations from 2021. To investigate these concerns, we conducted a comprehensive simulation study to examine potential safety implications arising from control transitions within mixed-autonomy traffic. The simulation results indicate adverse safety impacts due to increased safety-relevant interactions between vehicles caused by transitions of control in dynamic traffic flow conditions. Our findings also reveal that those effects could become stronger once string unstable ACC controllers are deployed as well.

Keywords: automated vehicles (AVs); Level 3 automation; mixed-autonomy traffic; surrogate safety measures (SSMs); take-over request (ToR); transition of control (ToC)

Citation: Alms, R.; Wagner, P. Control Transitions in Level 3 Automation: Safety Implications in Mixed-Autonomy Traffic. *Safety* **2024**, *10*, 1. https://doi.org/10.3390/safety10010001

Academic Editor: Raphael Grzebieta

Received: 2 November 2023
Revised: 5 December 2023
Accepted: 14 December 2023
Published: 19 December 2023

1. Introduction

With Level 3 automated driving systems at the verge of entering today's market [1,2], questions concerning traffic safety inevitably arise due to impending Level 3 disengagements that result in so-called 'transitions of control' (ToCs) [3]. Whereas contemporary research primarily emphasizes human factors in Level 3 conditional driving [4], an often neglected aspect is the potential macroscopic impact of ToCs on traffic safety caused by procedural effects in automated driving systems. As automation levels increase, so do the demands on human drivers in transition situations to ensure safe vehicle operation, since Level 3 operation permits drivers to divert their attention from present traffic situations (cf. SAE levels of automated driving [5]). Similarly, the technical prerequisites on the automation side must guarantee that handover situations are triggered within an appropriate time frame and under safe traffic conditions. UN Regulation 157 from 2021 initially approved Level 3 operation up to 60 km/h [6] for so-called 'Automated Lane Keeping Systems' (ALKS). Within these speed limitations, the first market-ready Level 3 systems have already been approved in Germany. Today's amended UNECE version from 2023 lays the regulatory foundation for certifying fully operational Level 3 ALKS even up to 130 km/h [7]. Prospectively, this opens up the opportunity for manufacturers to have their automated systems certified in stages, e.g., up to 80 km/h or, if the technical requirements and capabilities allow for it, for the full speed range up to 130 km/h. This raises the question of how such an Operational Design Domain (ODD) considering a speed limitation could impact traffic safety and whether interim approval up to a level like 60 km/h poses risks regarding mandatory control transitions. Therefore, we have conducted an extensive simulation-based safety analysis to examine basic, scalable effects, yielding the following contributions:

(1) We elaborate on our approach to model, parameterize, and simulate control transitions in a traffic scenario, encompassing the implications stipulated by the current UNECE regulations considering Level 3 systems that adhere to mandated speed limitations.
(2) We introduce a comparative method for evaluating surrogate safety measures (SSMs) based on histogram distributions, which can be visually represented in a heatmap.
(3) We present findings from three distinct use cases exploring fundamental ToC effects. The results from the main use case indicate ToC-induced safety impacts for high demand that is close to capacity limit in mixed-autonomy traffic with Level 3 shares $\geq$20%. The safety impacts assessed in the second use case can be attributed as negligible when taking an anticipatory property for the ToC preparation phase of Level 3 automation into consideration. The third use case shows increasingly detrimental traffic safety implications with ACC-related string instabilities in traffic flow.

Presently, comprehensive analyses concerning safety implications of ToCs are limited. We attribute this fact to deficiencies and constraints in existing ToC modeling. This study aims to provide more precise results with the help of a novel ToC model introduced by [8,9].

The rest of the paper is structured as follows: In Section 2 we elaborate on the control transition process, discuss simulation and modeling aspects, and review the latest research results assessing large-scale implications of ToCs. Moreover, we present selected SSMs that are relevant for our further investigation. Section 3 details our simulation experiment and the proposed methodology to assess the selected safety metrics. In Section 4, we present the simulation results and engage in a discussion of the findings. At last, in Section 5, we draw our conclusions from this study.

2. Literature Review

2.1. Transitions of Control

Control transitions, or transitions of control, delineate procedures that regulate the transfer of authority between the human driver and the automated vehicle (AV). The SAE taxonomy for driving automation systems [5] defines five different levels (cf. Figure 1); within such, ToCs constitute the transitions between those levels, typically from higher levels to zero or vice versa. In the case of upward ToCs, the driver cedes control to the AV within its predefined ODD. This transfer is facilitated through the activation of its automated driving systems, relinquishing the execution of longitudinal (acceleration/deceleration) and lateral (steering) driving tasks to those systems. For automation Levels 1–2, the human driver keeps monitoring the driving environment, whereas for Levels 3 up to 5, those monitoring tasks shift to the automation. Up to Level 3, the human driver poses as the fallback entity to ensure safe driving performance.

In instances for which the vehicle automation, due to a range of potential external or internal factors, cannot operate safely within its ODD anymore or is requested to disengage, a downward ToC is initiated. This prompts the AV to request the driver's intervention in order to continue manual control of the vehicle. The mechanism employed to inform the driver of the need to re-engage in primary driving responsibilities, which may involve auditory, visual, or haptic signals or combinations thereof from the vehicle automation, is termed a 'take-over request' (ToR). A successful downward ToC is considered complete once the driver has re-engaged and continues normal operation with fully restored situational awareness and driving skills. However, if the downward ToC fails, meaning the driver does not respond to the ToR within the specified lead time, the AV will perform a minimum risk maneuver (MRM) to come to a safe stop.

Level	Name	Narrative definition	DDT		DDT fallback	ODD
			Sustained lateral and longitudinal vehicle motion control	OEDR		
Driver performs part or all of the *DDT*						
0	No Driving Automation	The performance by the *driver* of the entire *DDT*, even when enhanced by *active safety systems*.	*Driver*	*Driver*	*Driver*	n/a
1	Driver Assistance	The *sustained* and *ODD*-specific execution by a *driving automation system* of either the *lateral* or the *longitudinal vehicle motion control* subtask of the DDT (but not both simultaneously) with the expectation that the *driver* performs the remainder of the *DDT*.	*Driver* and *System*	*Driver*	*Driver*	Limited
2	Partial Driving Automation	The *sustained* and *ODD*-specific execution by a *driving automation system* of both the *lateral* and *longitudinal vehicle motion control* subtasks of the *DDT* with the expectation that the *driver* completes the *OEDR* subtask and *supervises* the *driving automation system*.	**System**	*Driver*	*Driver*	Limited
ADS ("*System*") performs the entire *DDT* (while engaged)			*System*	**System**	*Fallback-ready user (becomes the driver during fallback)*	Limited
3	Conditional Driving Automation	The *sustained* and *ODD*-specific performance by an *ADS* of the entire DDT with the expectation that the *DDT fallback-ready user* is *receptive* to *ADS*-issued *requests to intervene*, as well as to *DDT performance-relevant system failures* in other *vehicle* systems, and will respond appropriately.				
4	High Driving Automation	The *sustained* and *ODD*-specific performance by an *ADS* of the entire *DDT* and *DDT fallback* without any expectation that a *user* will respond to a *request to intervene*.	*System*	*System*	**System**	Limited
5	Full Driving Automation	The *sustained* and unconditional (i.e., not *ODD*-specific) performance by an *ADS* of the entire *DDT* and *DDT fallback* without any expectation that a *user* will respond to a *request to intervene*.	*System*	*System*	*System*	**Unlimited**

Figure 1. SAE levels from [5] Table 1. The red rectangle highlights the fallback entity for Level 3 automation. Abbreviations: DDT = dynamic driving task; OEDR = object and event detection, recognition, classification, and response; ODD = operational design domain.

Particularly, the topic of takeover times in AVs in urgent situations has been researched in great detail over the past decade in human sciences, mostly in driving simulators aimed at specifying criteria for the appropriate design of human–machine interfaces for takeovers [10]. A comprehensive overview of such experimental studies is provided by [11]. Another motivational factor behind those research studies was to derive indicators of how to conceive takeover strategies in simulation models and automated driving systems with respect to lead times and post-ToC behavior. Recently, studies have focused on human responses to prototypical Level 3 automated driving systems in terms of physiological effects, risk acceptance, comfort level, trust, and various other aspects from the perspective of being in a passenger role during Level 3 operation [12–18]. An extensive literature review examining various influential factors on takeover performance is given by [19]. A notable approach of a generic multi-level framework for microscopic simulation by [20] incorporates human factors such as task demand, task capacity, and situational awareness. Based on this model, ref. [21] presented a detailed simulation study to investigate vehicle interactions, alluding to detrimental effects of accumulated ToCs in their analysis. However, most of the related research focuses on individual driving performance and local effects (collision avoidance, lateral and longitudinal safety, and risk tolerance) and is often conducted in driving simulators, small-sample-sized real-world tests, or sub-microscopic simulations. Thereby, the simplified state machine model developed by [8] for the microscopic traffic simulation SUMO [22], designed to efficiently capture potential disruptions in traffic caused by ToCs and to facilitate large-scale simulations, is further elucidated in the following paragraph.

Figure 2 presents an illustration of the model, with panel (a) depicting a generic velocity timeline a for successful transition and panel (b) showing a failed transition with a differing timeline. In both variants, a preparatory ToC phase is initiated after a ToR has been triggered, after which automated operation continues for a limited time span (referred to as the available lead time). Yet, the AV's automation system takes safety precautions during the preparatory ToC phase, i.e., increasing its headway, lane change avoidance, and acceleration abstinence, and continues operation until either re-engagement from the driver (successful ToC) or expiration of the available lead time (failed ToC), at which point it starts an MRM. For the latter case of a failed ToC, the model assumes a constant deceleration rate for the MRM, which can occur in the vehicle's current lane or the right-most lane (via a autonomous lane change maneuver) depending on traffic conditions. In the case of successful control transition, the post-ToC phase's driving performance is determined by a driver state model that considers perception errors for imperfect driving, as detailed in [8], which was partly adapted and based on action point models developed by [23–26].

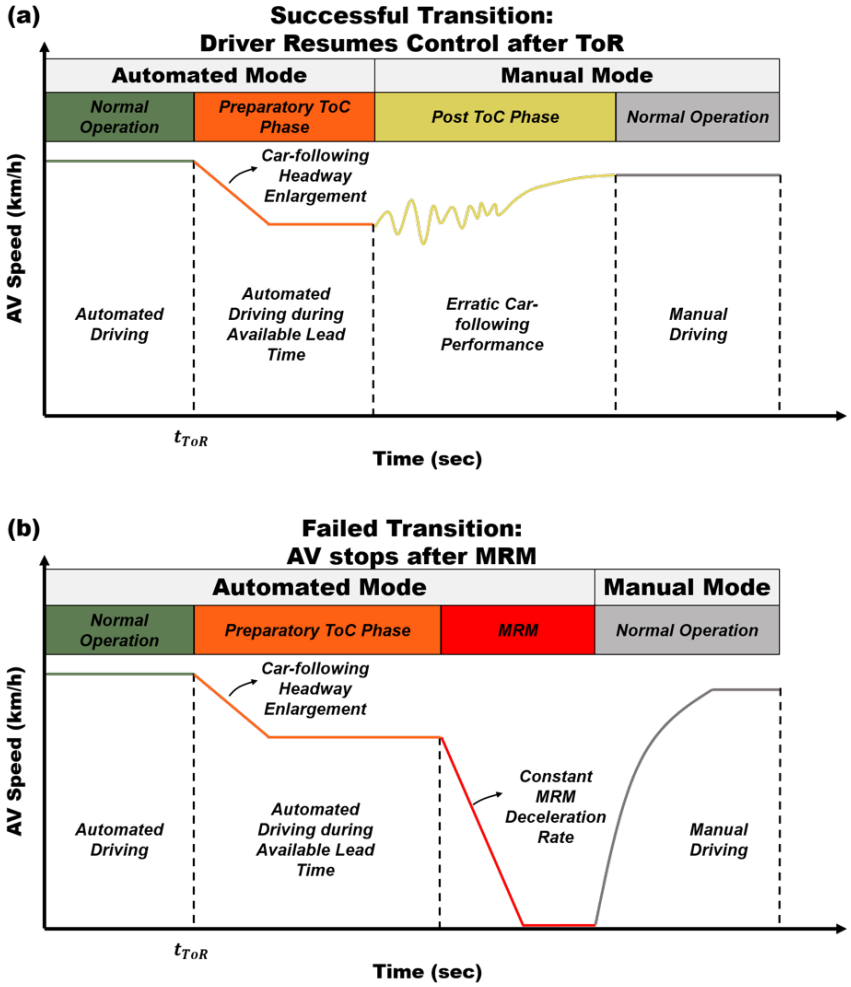

Figure 2. Illustration of ToC model mechanisms based on descriptions from [8]. Panel (**a**) depicts a generic velocity timeline for a successful transition; Panel (**b**) outlines a failed transition.

While previous research on ToCs so far has predominantly focused on individual consequences of failed transitions by analyzing crash incidents or disengagement reports [27–30], the examination of ToCs' large-scale effects on traffic has only gained attention in recent years. Among the first studies addressing complications induced by ToCs and exploring potential traffic management countermeasures is the work by [9]. This study identified adverse impacts on both traffic efficiency and safety. In a small simulation study emphasizing

traffic efficiency performance, Ref. [31] developed a simplified ToC model that restricts lane changes during control transitions while approaching a bottleneck. Although the study discussed safety-related aspects of ToCs, their analysis solely focused on traffic efficiency. Another very insightful study conducted by [32] assessed the potential impact of ToCs through a comprehensive simulation study for all of Japan. They analyzed variations in crash rates and highlighted the role of overconfidence or distrust in detrimental ToCs, which diminish the accident reduction effects they previously identified. The study generally observed an overall positive impact of vehicle automation on crash rates, particularly when the market penetration of AVs exceeded 50%. Furthermore, ref. [33] re-simulated specific motorway scenarios, such as 'cut-in' situations, based on accident data from the GIDAS dataset provided by the Federal Republic of Germany. This analysis focused on the severity of changes in these driving scenarios and derived potential benefits from automated driving functions by projecting the results to a national scale. In a simulation-based case study aimed to propose potential traffic management countermeasures to mitigate adverse impacts induced by control transitions [34], our earlier research focused on traffic efficiency performance, deferring an in-depth analysis of safety ramifications from ToCs. Therefore in Section 3, we introduce a new case study dedicated to investigating these safety effects in greater detail, and we share our findings in Section 4.

2.2. Surrogate Safety Measures

Surrogate safety measures are an important and helpful tool for evaluating traffic safety, particularly when crash events in data are rare or nonexistent, as in mixed-autonomy traffic scenarios. A good categorization, including mathematical definitions of relevant SSMs, can be found in a survey in [35]. Another comprehensive and quite deliberate literature review is given in [36], which details state-of-the-art SSMs in mixed-autonomy traffic research. Particularly, the authors critically discuss certain shortcomings of SSM-based safety assessments and point out inadequacies in simulation-based safety evaluations, which concentrate on automated driving without calibrated vehicle models. They also highlight underestimation of criticality and risk due to inadequacies in vehicle modeling, consideration of reaction times, and the interpretation of solitary SSMs with fixed but non-validated thresholds.

Typically, SSMs are categorized into two main classes: (I) SSMs for identifying individual conflicts and (II) SSM-based models for estimating crash risks or probabilities. Common sub-categories for (I) can be further divided into (i) time-based, (ii) deceleration-based, or (iii) energy-based, or combinations thereof. The time-based SSM we refer to hereafter in this work is the well-established Time to Collision (TTC) [37], which is defined as:

$$TTC = \begin{cases} \frac{d}{v_2 - v_1}, & \text{if } v_2 > v_1 \\ \infty, & \text{otherwise} \end{cases} \tag{1}$$

with d denoting the space gap and $v_2 - v_1$ being the speed difference between the leading and the following vehicles. Another representative deceleration-based SSM, denoted as the Deceleration Rate to Avoid Crash (DRAC), dates back to a concept introduced by [38] (although not yet fully formalized as a metric in that paper). It is defined as:

$$DRAC = \begin{cases} \frac{(v_2 - v_1)^2}{2d}, & \text{if } v_2 > v_1 \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

For an energy-based SSM, we reference [39] as an example, which presents an interesting, novel approach called 'extended delta-V', which estimates the crash severity of potential conflicts. Since this work does not delve deeply into SSM-based models or other related approaches, we refer the interested reader to [35,36] for more comprehensive information. However, we want to allude to a few notable papers that helped us develop a feasible approach for our own case study (see Section 3.3). In their study, ref. [40] employed an

approach for assessing the impact of connected automated vehicles on traffic safety by evaluating the distributions of TTCs rather than using a specific threshold for counting critical events. Even though their analysis is mostly descriptive, the histogram-based approach circumvents the utilization of non-validated criticality thresholds. Such so-called 'surrogate safety histograms' (SSHs) have also been used for completely different contexts, i.e., analyzing traffic safety based on vehicle trajectories for individual vehicle conflicts, and has proved to be helpful for comparing TTC results for different conditions [41,42].

As pointed out by [35,36], the absence of human perception–reaction times (PRT) or response times (RT) is a noted concern. In response to this issue, we introduce the Modified Deceleration Rate to Avoid Crash (MDRAC), which was developed by [43] as an enhancement over the traditional DRAC. The MDRAC takes into account a perception–reaction time (PRT) and is devised as follows:

$$MDRAC = \begin{cases} \frac{v_2 - v_1}{2(TTC - PRT)}, & \text{if } TTC > PRT, v_2 > v_1 \\ \infty, & \text{otherwise} \end{cases} \quad (3)$$

When [44] introduced their novel SSM, Deceleration Rate to Avoid a Crash using Constant Initial Acceleration (DCIA), they utilized the MDRAC for verification purposes. Their work inherently presented compelling data that underscored the superior sensitivity of the MDRAC compared to the traditional DRAC. (For their analysis, they used PRTs of 1.3 s and 2.02 s. We elaborate on our rationale for using a PRT of 1 s in Section 3.3.) The DCIA model was developed with the aim of accounting for the relative acceleration between two vehicles. Similar considerations were also explored by [45] but with a different analytical approach involving the design and introduction of a modified TTC, referred to as 'MTTC'. Both publications emphasize the limited expressiveness of the TTC when considered individually, motivating them to develop SSMs that incorporate relative decelerations. In a thoroughly performed study, ref. [46] validated bivariate SSM threshold pairs based on crash data, concluding that only a combination of conflict indicators can reliably estimate crashes. Their research also indicates that the DRAC, in particular, seems to systematically underestimate potential crash events.

Taking into account the previously mentioned considerations, we utilize two SSMs in our subsequent safety analysis: (1) the TTC and (2) the MDRAC. In a scenario-based simulation study, we investigate safety implications of ToCs, as specified in the following sections.

3. Simulation Experiment

3.1. Use Case Definition

Based on the principles set forth in UN Regulation 157 from 2021 [6] that led to first automated system approvals adhering to a speed limitation, we define a use case to investigate the impact of ToCs on traffic flows in mixed-autonomy driving conditions. Those regulations permit these approved Level 3 systems to currently operate up to 60 km/h in ODD compliant conditions, e.g., on unidirectional infrastructure such as highways. Therefore, on a two-lane highway, we split the speed limit into a sub-60 section, where AVs can drive in automated mode, and a downstream section with a higher limit of 100 km/h (cf. Figure 3a). Consequently, vehicles will accelerate to their desired or allowed speed (up to 100 km/h) when entering the latter section, requiring AVs to perform a ToC when their own current speed and also the perceived speeds of the surrounding vehicles exceeds the 60 km/h limit for Level 3 ODD operation.

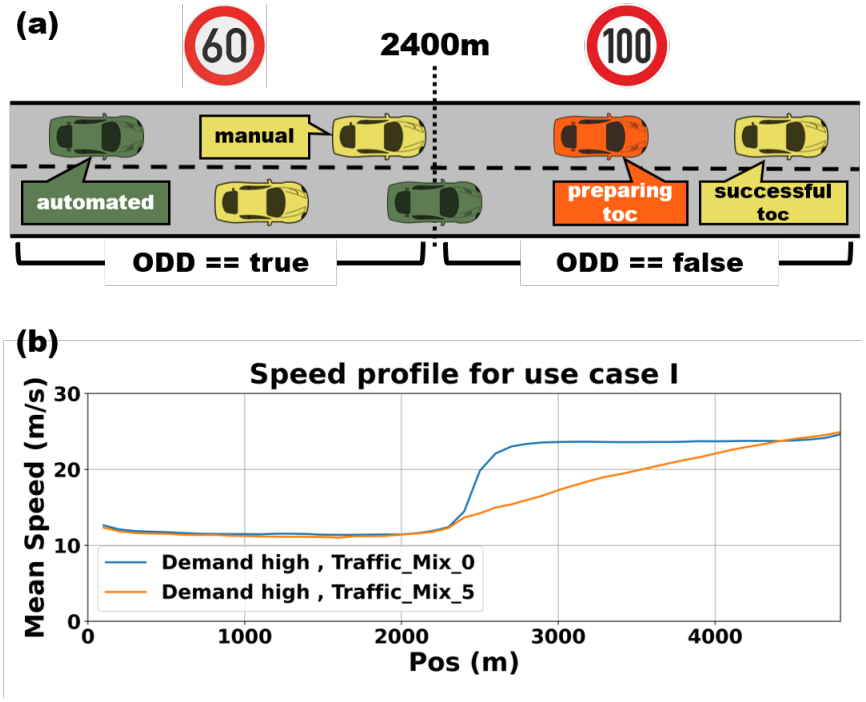

Figure 3. Panel (**a**) illustrates a schematic two-lane highway scenario where AVs are required to perform a ToC when exceeding 60 km/h. Panel (**b**) shows mean speed variants for this use case with different AV shares (Mix 0 versus Mix 5).

The speed profiles for two traffic flows in this use case are shown in Figure 3b and present the mean speeds for a traffic mix without AVs (cf. blue line, Mix 0) versus a traffic mix with a high AV share (cf. orange line, Mix 5). Comparing these two profiles, we see that due to consecutive ToCs, disruptions in the traffic flow delay the overall speed increase in the high-AV case.
Therefore, we define two scenarios:

- Base: AVs do not perform ToCs based on ODD limitations. Consequently, AVs continue to drive fully automated without speed limitations, accelerating up to 100 km/h in the second section of the road. All vehicles operate in their respective automation modes.
- Level 3: AVs operate within the ODD limitations. Therefore AVs are required to perform a ToC in the second section of the road, where surrounding traffic allows for driving at higher speeds beyond 60 km/h.

In this main use case, we investigate the principle safety implications caused by control transitions when the ToC model operates with the default parametrization: that is, emulating a preparatory ToC phase that actively increases a safe gap from its leading vehicle until the transition to the human driver is successful (cf. Section 3.2). UNECE regulations state that the automation *may* reduce the vehicle's speed during the transition to ensure safe operation. We discuss alternative approaches and potential countermeasures to avoid active deceleration by the vehicle automation in Section 4.2, but we also argue the technical requirements and difficulties that Level 3 automated vehicles face in such heterogeneous conditions.

3.2. Parametrization and Simulation Setup

We constructed traffic compositions with a broad range of vehicle shares, including AVs, manual vehicles (MVs), and heavy and light goods vehicles (HGVs and LGVs), as presented in Table 1. Each vehicle enters the simulation with its respective parametrization scheme as a manual or automated vehicle. LGVs and HGVs are considered to be manually driven trucks. AVs that do not perform a ToC in the Base scenario keep their ACC model

throughout the simulation. For the Level 3 scenario, AVs continue their operation as MVs after performing a ToC, which means they switch from the ACC to SUMO's default model. Table 2 summarizes the deployed models. Parametrization schemes for vehicles, lane changes, and ToC model(s) were adopted from [47]. The parameters used for defining the ToC models' preparatory phase are shown in Table 3. The simulations were performed with SUMO version 1.19.

Table 1. Traffic compositions with vehicle shares for six different mixes.

Traffic Mix	Vehicle Type			
	AV	**MV**	**HGV**	**LGV**
Mix 0	0%	85%	5%	10%
Mix 1	20%	65%	5%	10%
Mix 2	40%	45%	5%	10%
Mix 3	60%	25%	5%	10%
Mix 4	80%	5%	5%	10%
Mix 5	85%	0%	5%	10%

Table 2. Vehicle types in the simulation represented by SUMO model combinations. The Krauß model is SUMO's standard model. The ACC model is part of SUMO's car-following model selection.

Driving Mode	SUMO Model	Vehicle Type	
		MV/HGV/LGV	**AV**
Car-Following	Krauß	o	-
	ACC	-	o
Lane Change	Default	o	-
	Parameterized LC	-	o
Control Transition	ToC	-	o

Table 3. Parametrization of SUMO's ToC model.

Parameter	Description
`ogNewSpaceHeadway = 5.0 m`	Target additional space headway during the preparatory phase before a ToC
`ogNewTimeHeadway = 10.0 s`	Target time headway during the preparatory phase before a ToC
`ogChangeRate = 1.0`	Change rate of headway adaption during the preparatory phase before a ToC
`ogMaxDecel = 1.0 m/s`^{2}	Maximum deceleration rate due to headway adaption during the preparatory phase before a ToC
$t_{\text{lead}} = 10$ s	Available lead time for AVs

In regard to parameterizing t_{lead}, and also the other parameters in Table 3, it is important to note the lack of reliable information available from the manufacturers. In recent EC projects like L3Pilot or Hi-Drive, the manufacturers explicitly stated their restrictive data disclosure: referring to a competitive customer market. Recent technical reports from the Hi-Drive project [48,49] suggest a specification range from 5–20 s but without further details on when ToRs will be triggered. Moreover, in a press release, [1] stated 10 s as a lead time for their Drive Pilot system.

In order to compare both scenarios accurately, based on volume/capacity ratios, we determined the capacities individually for each scenario and each traffic mix. This was

deemed necessary since the six distinct traffic compositions defined in Table 1 entail different capacities due to differing lane utilization caused by the vehicle model's longitudinal and lateral driving behavior (including capacity-diminishing effects caused by ToCs). To identify the respective capacity for each traffic mix, we ran simulations with increasing demands until a surge in the vehicle insertion backlog could be detected, which means that SUMO cannot insert more vehicles per time step in a fully populated network (SUMO delays vehicle departures to ensure safe gaps between consecutive vehicles if the minimum gap would be violated). Table 4 following summarizes the measured lane capacities for both scenarios.

Table 4. Measured lane capacity c $(veh/hour)$ per traffic mix in SUMO for both scenarios .

	Traffic Mix					
	Mix 0	**Mix 1**	**Mix 2**	**Mix 3**	**Mix 4**	**Mix 5**
Base	1700	1600	1550	1500	1450	1450
Level 3	1700	1500	1350	1250	1230	1170

Vehicles randomly enter a 5000 m, two-lane road network with a Poissonian distribution, while the demand is determined for a volume/capacity ratio range from [25–100%] per parameter combination traffic mix/percentage. To represent the vehicle types (cf. Table 2) with a heterogeneous fleet, we created random distributions for each type, wherein SUMO randomly picks vehicles for insertion from a type-specific set of 1000 differently parameterized vehicles each time step.

For the Level 3 scenario, in order to emulate the necessary return of vehicle control from the automation to the human-driver, ToRs are triggered for AVs through the SUMO API TraCI in the second road section. ToRs are issued based on the current ego vehicle speed and the average speed of the surrounding road section. A simplified abstraction of the mechanism is shown in the following pseudo-code:

```
while simulation = true do
  monitor speedLevel
  add AV to pendingVehiclesList
  for AV in pendingVehiclesList: do
    if (AV.speed && speedLevel > 16.7 m/s) then
      issue ToR
    else if AV.pos ≥ endOfSection then
      issue ToR
    end if
  end for
end while
```

For the Base scenario, no alteration to vehicle behavior through TraCI is necessary. To model individual human driver takeover times to ToRs, a Gaussian probability distribution was created as in [47] with an expectation of 7 s and a variance of 2.5. This distribution results in about 10% short MRMs: i.e., driver reaction times exceeding the available lead time of 10 s but only with durations of less than 3 s. MRMs with longer durations up to causing the standstill of a vehicle, thereby blocking a full lane, are not considered in this study. A concept of how to handle long MRMs with infrastructure assistance deploying V2X messages can be found in [9,50].

Takeover time distributions are often analyzed in controlled driving simulator studies such as, e.g., [51], but clear evidence of the distribution pattern (normal, skewed, platykurtic,...) in regard to Level 3 related ToCs remains inconclusive, which is why we refer to the stated normal distribution. The mean takeover time is set seemingly high compared to takeover times often referenced in studies for urgent ToCs with rather short takeover time budgets. A range of 7–8 s as a common time budget for supposedly urgent ToCs with a respective range of approximately 2–3.5 s for takeover times is stated by [19] in their review.

Similar ranges were previously identified in a survey by [52], but their own research also focused on non-urgent, normal ToCs, which showed notably increased takeover times with secondary task involvement (median: 6 s). Since we consider the ToCs mandated in our described use case as non-urgent, planned ToCs, we regard the referenced mean takeover time of 7 s as adequate for our parametrization.

All parameter combinations run with 10 different seeds for at least 1 h simulated time. In total, we conducted 6 mixes × 6 ratios × 10 seeds × 2 scenarios = 720 simulations for the main evaluation (plus an additional 1440 simulations for the two extra use cases: cf. Section 4).

3.3. Method for Assessing SSMs

We assess the safety ramifications with two SSMs: TTC (time-based) and MDRAC (deceleration-based). Typically, theses metrics are count based: meaning the number of events within a specific time period is accumulated. Therefore a dedicated threshold needs to be specified to define an event as critical. These thresholds, which differentiate events between critical and non-critical, have a wide range in the literature and also depend on the context of real-world, driving simulator or simulated traffic data. Based on crash evaluations from naturalistic driving studies (NDS), ref. [53] suggests a TTC range 0.7–1.4 s, whereas [54] states a range of 2.6–3 s from driving simulator tests. Reference [55] summarizes its findings as a range of 1.5–5 s for TTCs. For MDRACs, usually the same range from 3.0–3.4 m/s^2 as for the DRAC is utilized [44,46]. (A basic evaluation with what we found to be two illustrative thresholds each (i.e., TTC with 1.75 s and 3.0 s; MDRAC with 3.0 m/s^2 and 3.4 m/s^2) is shown in Section 4).

On the other hand, we note the quite limited expressiveness of such a threshold-based evaluation since it is bound to the validity of the specified thresholds that define the criticality. For SUMO, no validated set of thresholds exists for any of the established safety metrics since those would also have to be calibrated for each car-following model individually. Due to no available real-world data with Level 3 automated vehicles on a large scale, we see no feasible option to find a valid set of criticality thresholds in SUMO, at least for our scenarios.

Thus, we offer a complementary analysis with the selected SSMs that aims to provide better comparability between the two defined scenarios by exposing differences more robustly than a count-based evaluation could, at least for our use case. Figure 4 exemplarily shows the histograms and density distributions for TTC and MDRAC data of one simulation set (i.e., 10 seeds), comparing the two defined scenarios: Base vs. Level 3. In this case, the histograms illustrate rather substantial differences between the distributions (orange vs. blue) for both metrics. Consequently, distinguishing between critical vs. non-critical events by appointing a distinct critical threshold might be obvious and robust for identifying safety issues. On the contrary, when data points are sparse or roughly similar, noticing safety issues from the data might be more inconclusive because even small differences in an appointed criticality threshold may have great impacts on the total number of count-based events. (We consider it impractical to include all 72 histograms from our primary simulation case. The presented heatmaps encapsulate these histogram data.)

The proposed steps to evaluate a full dataset per SSM are listed as follows:

1. Include all data points within a range of interest of the SSM (i.e., TTC events between 0–8 s and MDRAC events > 0.2 m/s^2) and calculate the surface integral of the histogram. This results in one condensed value for the total number of TTC or MDRAC events, respectively, for a full simulation set
2. In order to compare both scenarios against each other, calculate the difference between the surface areas derived from the histograms: Total = Level 3 − Base.
3. Draw a colored heatmap for all parameter combinations for each metric. Also, normalize the data for better comparability to other (alternative) scenarios.

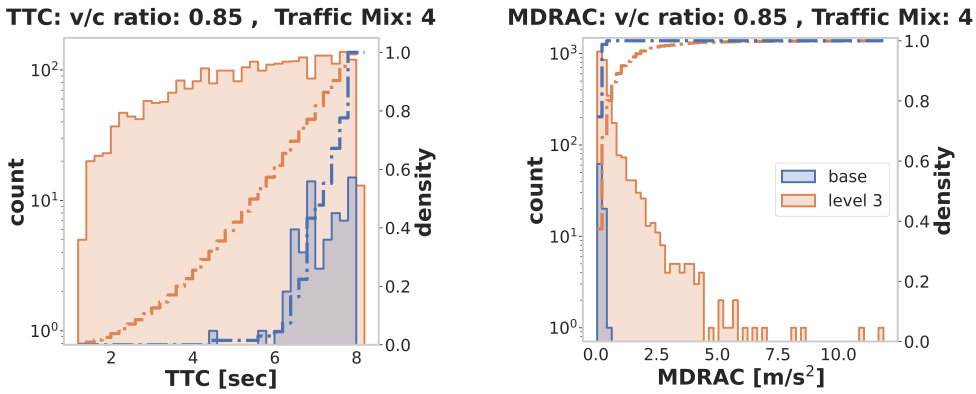

(**a**) TTC (**b**) MDRAC

Figure 4. SSMs for main use case: parameter combination traffic Mix 4, v/c ratio 85%. Panels (**a**,**b**) present histograms for TTC and MDRAC on a log scale (left y-axis) comparing Base (orange) vs. Level 3 (blue); cumulative densities are represented by dash-dotted lines (right y-axis).

All events are weighted equally when calculating the integral. Inherently, events with higher TTC values occur more often than those with lower TTC values in traffic flows; thus, the surface value holds greater emphasis on more frequent but potentially less critical events (and vice versa for MDRAC events). Additional tests with weighted results showed negligible impact on the final heatmap visualizations, supporting the robustness of our approach when sufficient data are available.

Note, that for each traffic mix defined in Table 1, there is a corresponding capacity limit referenced in Table 4. As described before, these different capacities entail that in order to isolate the impact of ToCs within each traffic mix, all parameter combinations of traffic mixes and v/c ratios have to be simulated with and without the occurrence of ToCs. This is why scenarios Base and Level 3 were defined in Section 3.1 and why calculating the difference between those respective results, as denoted in Point 2, is deemed necessary. To further highlight the safety impact of ToCs, we include traffic Mix 0 in all following representations of the results as a reference traffic share without AVs (and consequently, no ToCs).

Additionally, Figure 4 shows cumulative density distributions on the the second y-axis (see dashed lines). Those distributions can also be used for non-parametric statistical tests, such as Kolmogorov–Smirnov [56] or Cucconi [57]. Although both tests were carried out on the result data to test for statistical significance, we do not present these explicitly, since non-parametric tests usually are sensitive to the sample size. Note, that we ran those tests for accumulated results (high sample size) as well as separately on each seed (relatively small sample size) for all parameter combinations of mixes and v/c ratios. In most of the cases, the tests showed highly significant differences that are in line with the presented heatmap results shown in Section 4. Nevertheless, we refrain from including these statistics since we cannot determine a suitable sample size due to the nature of our simulated data from the scenario at hand. Moreover, for evaluating the MDRAC, we implemented the necessary code for the SSM device in SUMO (see https://github.com/eclipse-sumo/sumo/pull/13352, accessed on 1 November 2023) and parameterize it with a PRT value of 1 s. For simplicity, we keep this value constant throughout the entire evaluation. As [58] stated in his extensive analysis, there is no exact value to estimate the reaction time for human drivers to brake, since it depends on the level of expectancy. The author provides a range from 0.75 s up to 1.5 s for mean response times. Note also, that in this context, this means we only assume a PRT to analyze the SSM. SUMO does not model a PRT in its car-following models per se. The parameters `stepLength` and `actionStepLength` for all simulations are set to 0.1 s; accordingly, the vehicle models react in each time step. Even though decoupling theses two parameters can emulate aspects of a PRT, it would still not resemble a true model of a

perception–reaction loop. This topic is worth a comprehensive study itself, but this is not the main focus of this work.

4. Results and Discussion

In the following section, we analyze and discuss the main results from the simulation experiments (cf. Section 4.1 as well as two alternative use cases (cf. Sections 4.2 and 4.3)—which examine different premises for simulating Level 3 automation—in order to disclose further potentially beneficial and detrimental safety ramifications.

4.1. Safety Analysis

Tables 5 and 6 present the count-based evaluation for critical *TTCs/hour* and *MDRACs/hour* with two different thresholds for each metric. Both tables only show values for the Level 3 scenario since no critical events could be detected for the Base scenario.

Focusing on Table 5, we observe that for either threshold, with very low demands i.e., v/c ratios below 55% and without automated vehicles in the traffic composition, i.e., Mix 0, zero or only very few events per hour can be noted. With increasing demand (>55%), critical events occur more frequently, even with lower automation shares for Mix 1 and Mix 2. In comparison, we notice obvious discrepancies in the number of critical events depending on the threshold (i.e., 1.75 s and 3.0 s). Nevertheless, for both thresholds, similar trends for mixes $\geq$ Mix 2 and demands closer to capacity limits $\geq$ 70% can be identified. For the maximum values, the critical event rate increases from zero to either 7.3/h or up to 44.8/h depending on the threshold. Considering MDRACs in Table 6, we observe a less distinct discrepancy between the events for both MDRAC thresholds. Comparable trends to those stated for TTCs can be observed, with absolute values overall being closer to their corresponding TTC value pairs for the critical TTC threshold $\leq$ 1.75 from Table 5. Maximum values for increased critical event rates are 7.1/h and 8.9/h, respectively.

Table 5. Comparison of critical TTCs/h for Level 3 scenario with v/c ratio range from [25–100%] for traffic Mixes 0–6. The critical thresholds chosen for the measurements are: TTC$\leq$ 1.75 s and TTC$\leq$ 3.00 s. Note that for all parameter combinations in the Base scenario, no critical events were detected.

		Traffic Mix											
		Mix 0		**Mix 1**		**Mix 2**		**Mix 3**		**Mix 4**		**Mix 5**	
		≥1.75	**≥3.00**	**≥1.75**	**≥3.00**	**≥1.75**	**≥3.00**	**≥1.75**	**≥3.00**	**≥1.75**	**≥3.00**	**≥1.75**	**≥3.00**
v/c ratio [%]	100	0	0	6.5	29.1	6.4	36.6	6.8	37.6	7.3	44.8	7.1	39.7
	85	0	0	3.5	18.4	4.1	20.5	5.0	25.9	4.5	30.2	2.6	23.2
	70	0	0	2.7	9.6	3.2	14.8	3.1	14.8	4.1	22.3	2.7	16.8
	55	0	0	1.4	4.9	1.6	7.9	1.7	8.0	2.3	10.5	0.9	6.9
	40	0	0	0.2	1.9	0.4	1.9	0.6	3.7	0.4	3.4	0.7	3.1
	25	0	0	0	0.2	0	0.2	0.2	0.8	0.1	0.5	0.1	0.7

Correspondingly, Figure 5 visualizes these tabular data exemplarily for critical thresholds $TTC \leq 1.75$ s and $MDRAC \geq 3.4$ m/s^2 in a normalized heatmap. We observe rather dissimilar coloring comparing both heatmaps, which makes it hard to clearly identity trends, particularly across traffic composition. For that reason, as proposed in Section 3.3, we present additional heatmaps that are detached from specific criticality thresholds and that aim to constitute more coherent results across both metrics. The heatmaps in Figure 6 provide more subtle insight into the distribution for increased TTC and MDRAC rates. Firstly, we observe reasonably matching coloring between both heatmaps: TTC vs. MDRAC. Second, to check for plausibility, we would expect for Mix 0 (which has no AVs in the fleet mix) to show no differences between the results, which the minimum light-green coloring in the heatmaps confirms. Similar coloring is to be found for v/c ratios with lower demands, i.e., 25–40%, displaying no noteworthy differences. Third, maximum

values appear on capacity limit in Mixes 2 and 4, as opposed to Figure 6, which shows Mixes 4 and 5 at maximum. Fourth, slightly more pronounced highlighting of MDRAC events—particularly for Mix 4/85% and Mix 5/100% but also seen as moderately darker colors for medium-demand levels—can be observed. This accentuation compared to the TTC coloring might hint at stronger sensitivity for the MDRAC versus the TTC.

Table 6. Comparison of critical *MDRACs/hour* for Level 3 scenario with v/c ratio range from [25–100%] for traffic Mixes 0–6. The critical thresholds chosen for the measurements are: $MDRAC \geq 3.4$ m/s^2 and $MDRAC \geq 3.00$ m/s^2. Note that for all parameter combinations in the Base scenario, no critical events were detected.

		Traffic Mix											
		Mix 0		**Mix 1**		**Mix 2**		**Mix 3**		**Mix 4**		**Mix 5**	
		≥3.40	**≥3.00**	**≥3.40**	**≥3.00**	**≥3.40**	**≥3.00**	**≥3.40**	**≥3.00**	**≥3.40**	**≥3.00**	**≥3.40**	**≥3.00**
v/c ratio [%]	100	0	0	3.3	5.1	3.8	5.7	6.5	7.8	6.0	8.9	7.1	8.9
	85	0	0	2.5	3.4	2.6	3.4	4.0	4.8	3.7	4.7	2.8	3.6
	70	0	0	1.7	2.2	2.1	2.8	3.6	4.1	4.4	5.3	2.6	3.2
	55	0	0	1.2	1.5	0.6	1.0	1.7	1.9	2.1	2.8	1.3	1.7
	40	0	0	0.2	0.4	0.2	0.3	1.1	1.5	0.7	0.8	1.0	1.2
	25	0	0	0.1	0.1	0	0	0.1	0.1	0.1	0.1	0.3	0.3

Nevertheless, the results seem to indicate adverse safety implications close to the maximum v/c ratio starting with traffic Mix 1 for both SSMs. Overall, these heatmaps showcase, from our point of view, more consistent trends regarding safety impacts of ToCs across the stated parameter combinations of demand/traffic mix than the threshold-based results alone could do (cf. Tables 5 and 6). (We conducted additional simulations with a similar use case that mandates a different speed limitation (80 km/h) for a Level 3 automated system. In this case, the network speed limit allowed speeds up to 130 km/h in the second road section. The overall trends shown in Figure 6 remain about the same for such a use case, even though the absolute numbers are lower due to considerably different capacities, which also lead to varying interactions between vehicles).

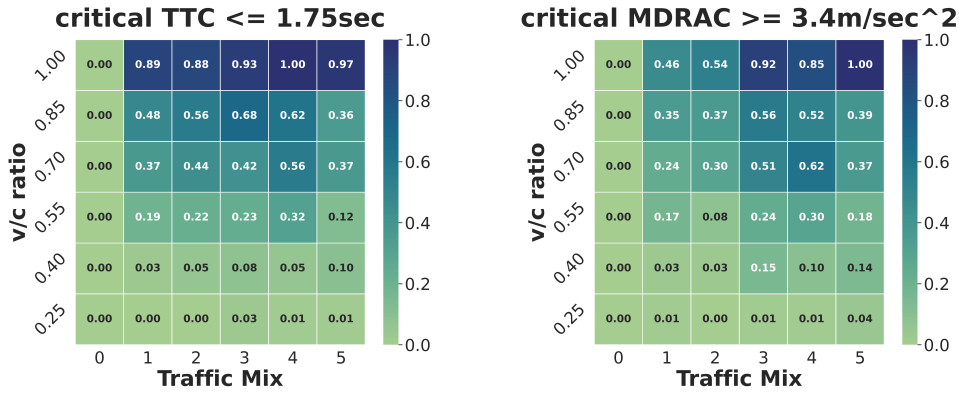

(**a**) TTC (**b**) MDRAC

Figure 5. Heatmaps for SSMs TTC (panel (**a**)) and MDRAC (panel (**b**)) across v/c ratios and traffic mixes for critical thresholds $TTC \leq 1.75$ s and $MDRAC \geq 3.4$ m/s^2 and normalized to respective maximum values.

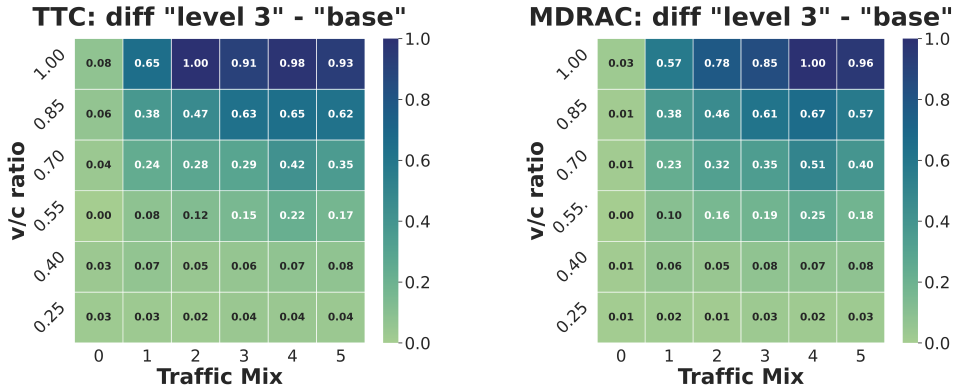

(**a**) TTC (**b**) MDRAC

Figure 6. Heatmaps for SSMs TTC (panel (**a**)) and MDRAC (panel (**b**)) across v/c ratios and traffic mixes and normalized to respective maximum of the difference between the surface integrals of the histograms.

4.2. Open Gap Prediction

As previously discussed in Section 3.1, we acknowledged that the ToC model in SUMO is intentionally designed and parameterized to actively increase the gap to its leading vehicle during the preparatory ToC phase. This design objective is aimed at ensuring a safe gap when the human driver takes back control. Considering that in this use case the automated Level 3 operation is only allowed to go up to 60 km/h before a ToC is mandated, numerous hypothetical factors could be listed on why traffic is moving below this threshold, e.g., speed limits, construction sites, accidents, weather conditions, infrastructure limitations, or high traffic density, to name a few.

Hence, when the factors causing the traffic slowdown subside and overall traffic conditions permit higher driving speeds, a Level 3 automated system could potentially anticipate the acceleration of the vehicle ahead. As a result, this assumption about surrounding traffic's acceleration would enable the AV, during a ToC preparation phase, to wait for the gap to naturally increase by maintaining its current speed—without the need for active deceleration—while the preceding vehicle accelerates. This mechanism may temporarily deviate from the minimum safe gap assumption of the ACC model for a few seconds until the gap becomes sufficiently large. Assuming that a Level 3 AV can accurately predict the behavior of the leading vehicle over a short time span, this approach might be considered an acceptable risk. It is important to note that current ACC systems permit users to manually set a desired minimum gap below a safe distance level. However, this is because human drivers are expected to continuously monitor the traffic situation and be prepared to regain control immediately (cf. SAE levels 1 and 2, Figure 1). In contrast, Level 3 automated systems must be capable of autonomously monitoring surrounding traffic conditions and reliably predicting a safe driving status within their ODD conditions. (A typical approach often discussed in research to have better predictability is the deployment of vehicle-to-vehicle communication (V2V), which is not in the scope of this work. Although, SUMO provides a respective CACC vehicle model, we keep the ACC model throughout our simulations for coherent comparability across the use cases.)

Therefore, to investigate potential benefits to the traffic flow presuming the discussed advanced Level 3 capabilities, we define another use case for which we perform reruns of the previous simulations but with different parametrization for the ToC model's preparation phase. The scenario is denoted as *OG* for the following evaluation. It is important to clarify that we do not explicitly model the advanced predictive capabilities of an AV in SUMO directly. Instead, we emulate the AV's predictive mechanism by adjusting the ToC model's parametrization. Specifically, the maximum deceleration during the headway adaption is set to `ogMaxDecel` $= 0\,\mathrm{m/s^2}$, thereby preventing any deceleration during the preparatory

ToC phase. This approach aligns with our assessment since none of our vehicle models or individual vehicles in the scenario configurations are determined to engage in unpredictable or abrupt braking maneuvers.

In line with our previous evaluation, Figure 7 presents heatmaps for both safety metrics. In order to compare the new open gap effects with the main use case results, we kept the coloring scale from Figure 6. Consequently, when observing the uniform light-green coloring of both heatmaps, it becomes evident that all detrimental control transition effects vanish completely compared to the main use case. Although this use case here represents a simplified version of the discussed predictive mechanism, the results highlight the potential of advanced Level 3 automation that not only accounts for individual safety but also the effects on overall traffic as long as it adheres to the criteria set by UNECE regulations to ensure safe control transitions. In our assessment, this requirement poses a significant technical challenge from an engineering and manufacturer perspective, making it a crucial topic for future research and real-world testing in the context of high-level automated driving.

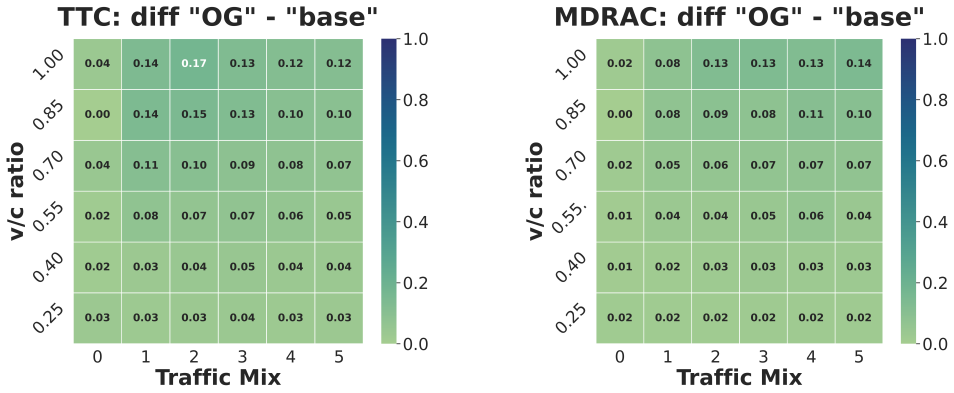

(**a**) TTC (**b**) MDRAC

Figure 7. Heatmaps for SSMs TTC (panel (**a**)) and MDRAC (panel (**b**)) across v/c ratios and traffic mixes: use case without deceleration during the preparatory ToC phase, which actively opens gap to leading vehicle; normalized to respective maximum of the difference between the surface integrals of the histograms from Figure 6.

4.3. String Stability

One aspect about deploying the ACC model as a proxy for Level 3 automated driving in SUMO that has not been discussed yet is the issue of string stability vs. instability. Although [59] discusses the possibility of parameterizing their originally proposed ACC controller for string stability, they presented calibrated gain factors ($k_1 = 0.23$ s^{-2} and $k_2 = 0.07$ s^{-1}) based on real-world experimental data. Notably, these gain factors resulted in unstable ACC string behavior (cf. [59] Figure 12). As shown in Figure 8, we conducted simulations in SUMO to replicate the driving maneuvers from the experiments presented in [59]. The speed oscillations of the subsequent vehicles clearly illustrate the string instabilities for the ACC model when parameterized with the previously stated control gains.

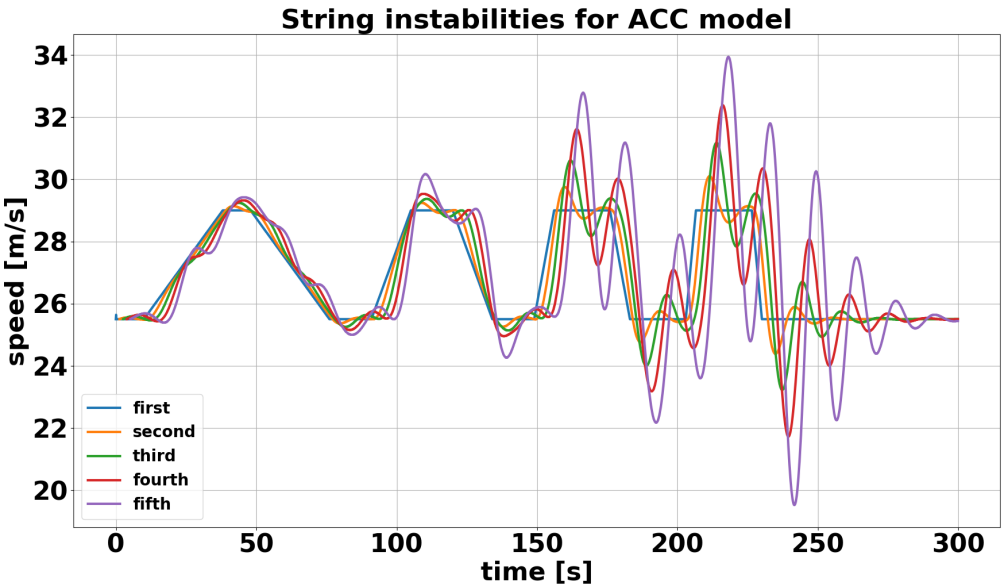

Figure 8. String instabilities of the ACC model as presented in ([59] Figure 12), re-simulated with SUMO (modes (iii) and (iv) deactivated).

In order to ensure a stable controller, [60] presented a modified ACC model with additional driving modes: adding a gain factor $k_0 = 0.4\ s^{-1}$ for a speed control mode plus a new mode called 'gap-closing', which tunes the gain factors k_1 and k_2 to $0.04\ s^{-2}$ and $0.8\ s^{-1}$. SUMO's implementation of the ACC model is based on this paper, but the model is extended with another fourth mode based on [47]. This mode is named the 'collision-avoidance' mode, and it tunes the gain factors to $k_1 = 0.23\ s^{-2}$ and $k_2 = 0.8\ s^{-1}$. All these modes and tuned gain factors aim to ensure that the ACC vehicle can brake hard enough to avoid collisions and therefore provide a crash-free simulation while resulting in a string-stable ACC parametrization.

However, other studies such as as [61,62] are based on experimental campaigns with ACC-equipped, commercial vehicles currently deployed to the customer market and demonstrate string instabilities in vehicle platoons. As the traffic compositions utilized in this work contain increasing AV shares, we would also expect a ToC-related safety impact in the presence of string unstable behavior. Therefore, in order to investigate potential detrimental ToC-related effects in string unstable traffic flow conditions, we conducted reruns of the main use case, deploying the parametrization scheme for string unstable ACC behavior, as depicted in Figure 8. Accordingly, Table 7 shows the updated measured lane capacities for this parametrization. (Considering the numerous collisions observed in these simulations, it is likely that the actual capacities are lower, but nonetheless, we assess those numbers as an adequate approximation. The differences in capacity between both scenarios cannot be discerned without more compartmentalized measurement because the strong string instability effects of the ACC model obscure the underlying ToC effects.)

Table 7. Measured lane capacity c $(veh/hour)$ per traffic mix in SUMO for both scenarios with string unstable ACC parametrization.

	Traffic Mix					
	Mix 0	**Mix 1**	**Mix 2**	**Mix 3**	**Mix 4**	**Mix 5**
Base	1700	1600	1550	1500	1450	1450
Level 3	1700	1600	1550	1500	1450	1450

Consistent with our proposed evaluation method, we created heatmaps for both SSMs (cf. Figure 9). For comparison, the coloring is normalized to the notably lower scale of

Figure 6. The black tiles highlight the data points (which are, as a reminder, the difference between the surface integrals of the histograms) that exceed the data from the main use case discussed in Section 4.1. We observe a significant spike in TTC and MDRAC events with increasing demand and AV shares: essentially for all mixes from Mix 1 to Mix 5 and v/c ratios greater than 40%. Although there is a notable difference between the absolute values between both SSMs when comparing these heatmaps, the overall trend towards higher demands and AV shares persist. Again, as noted before in Section 4.1, the MDRAC metric seems to be more sensitive to escalating safety events. Overall, we think this use case vividly demonstrates potential adverse ToC effects for unstable ACC controllers that may need to be carefully addressed in mixed-autonomy scenarios presuming Level 3 systems will reach significant market penetration under the considered regulations that mandate speed limitations.

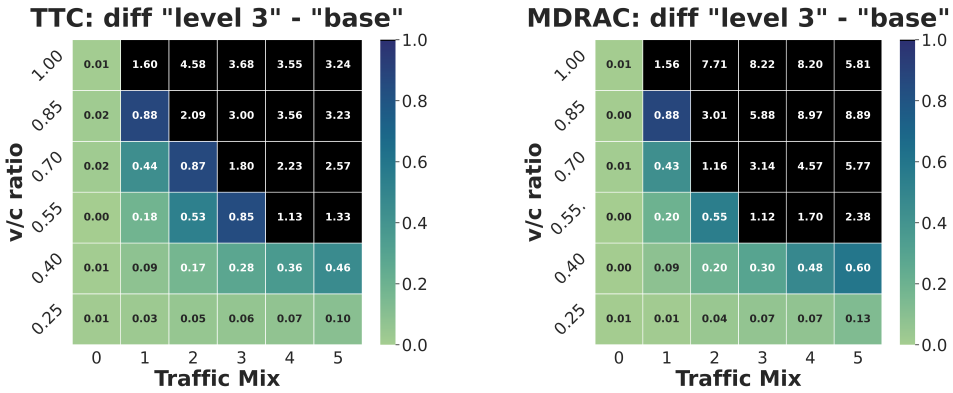

(**a**) TTC (**b**) MDRAC

Figure 9. Heatmaps for SSMs TTC (panel (**a**)) and MDRAC (panel (**b**)) across v/c ratios and traffic mixes: use case with string unstable ACC model parametrization; normalized to respective maximum of the difference between the surface integrals of the histograms from Figure 6.

4.4. Limitations of the Study

However, we want discuss various limitations of the study concerning the interpretation of the presented results.

4.4.1. Validation

As previously mentioned, SSM-based evaluations in traffic simulations often suffer from insufficient calibration to real-world data. Likewise, traditional count-based SSM analyses that employ specific thresholds to distinguish critical from non-critical events lack proper validation within microscopic traffic simulations like SUMO. Considering the limited deployment of Level 3 automated driving systems in the market, data are scarce due to manufacturers in a competitive customer market tending to withhold detailed information about automated driving operations, including parametrization schemes for ToC procedures. Unfortunately, this issue appears to be difficult to address in the short term. Yet, a promising option for better calibration, at least for SUMO's current ACC model, is the database OpenACC described by [63], which provides large experimental datasets on state-of-the-art Level 2 systems.

4.4.2. Safety Metrics

Our analysis relies on SSMs that are suitable for identifying safe longitudinal distances or time gaps, i.e., the TTC and the MDRAC. These SSMs have their limitations, particularly in terms of expressiveness, as they assume constant speed or deceleration. The TTC, e.g., is often inherited in modified metrics to account for such deficiencies. On the other hand, the MDRAC incorporates the PRT to address human response delays. However, we lack

information about a sensitivity analysis for a feasible distribution of PRTs calibrated for a microscopic simulation like SUMO. Another aspect that we have not yet discussed is the interdependency between longitudinal and lateral traffic effects. Increased longitudinal dynamics resulting from ToCs could potentially lead to increased lateral maneuvering by following vehicles. Measuring diverging lane change rates may provide further insights into ToC-related safety effects.

4.4.3. Quantification

By evaluating safety effects based on SSMs as TTC or MDRAC, we find it hard to quantify the potential increase or decrease in traffic safety compared to other, more conclusive indicators such as, e.g., definite collisions. This is again due to the limited validity of such safety measures in simulations without calibration to real-world data. Even though our findings indicate safety-related effects of ToCs, we cannot say anything about the severity of these impacts without assessing the results further, e.g., with the help of an energy-based SSM as in the previously mentioned *extended delta-V* [39]. Another limiting aspect we want to mention might be that the selected SSMs do not account for possible speed dependencies. Their expressiveness might also depend on individual vehicle types (cf. MVs, AVs, HGVs, and LGVs), which we did not factor in.

4.4.4. Modeling

We previously discussed certain aspects of ACC parametrization in relation to evaluating ToC implications in Section 4.3. In a broader context, we do not consider the ACC model to be an ideal proxy for simulating Level 3 automated systems. This is because the ACC model was originally developed to emulate a Level 2 adaptive cruise controller. Level 3 systems are expected to exhibit different gap control and gap-closing behavior due to distinct safety criteria outlined in the SAE levels for automated driving (cf. OEDRs, Figure 1 and [5]). A driver model worth considering might be the fuzzy logic model called 'FSM' that has been successfully proposed by [64] to be included to the current Regulation 157 for evaluating safety performance of the relevant critical scenarios. Regarding the parametrization of the ToC model, we acknowledge that the current model does not account for situationally variable lead times. Such variability could be attributed to a performant Level 3 system that can adjust its lead time based on current traffic conditions. Also, the distribution of human takeover times in Level 3 systems is an ongoing research topic, and our approach could be improved with better tailoring of the distribution form to suitable experimental data.

5. Conclusions

With consideration of an operational speed limitation up to 60 km/h as stated for today's approved Level 3 automated driving systems, we conducted a comprehensive simulation study comprising three different use cases to investigate potential safety implications of ToCs. Each use case demonstrates fundamental mechanisms and effects of control transitions in traffic flow, with the premise being that ToCs are mandatory in those circumstances. In our simulations, distinct parametrizations regarding ToC and ACC modeling serve as proxies for various considerations of what we believe to be potentially influential factors in future Level 3 capabilities. This research aims to address safety-critical aspects of Level 3 automated driving, including control transitions with the help of SSMs (TTC and MDRAC) in a field with limited large-scale data from real-world testing due to negligible market penetration of Level 3 systems thus far. The results can be summarized as follows.

(1) The main use case presents traditional threshold-based SSM results that may indicate adverse safety effects of ToCs for high demands and increasing AV shares, but clear trends remain inconclusive since both metrics diverge across the parameter combinations of v/c ratio and traffic mix. Our proposed analysis based on histogram data provides more

coherent trends for both SSMs that hint at potentially detrimental ToC impacts starting with AV shares of 20% and v/c ratios $\geq$ 70%.

(2) When considering advanced capabilities of Level 3 systems to perceive and predict the disintegration of continuous operation within an ODD under accelerating traffic conditions, we emulated these capabilities with an anticipatory ToC preparation phase that does not actively increase its gap distance to a preceding vehicle. Our results show that adverse safety effects completely dissipate under such a premise.

(3) Considering that nowadays deployed ACC controllers in automated driving systems might induce string instabilities in traffic flow for higher market penetrations rates, we reran simulations with a string unstable parametrization scheme for AVs. The results show significant spikes in TTC and MDRAC events for mixes $\geq$ Mix 1 and v/c ratios $\geq$ 40%, indicating negative safety impacts due to ToCs in heterogeneous traffic conditions.

(4) In Section 4.4, we address several limitations of our study. The primary limitations stem from the lack of validation due to limited real-world data and calibrated simulations, as well as potential shortcomings in vehicle and ToC modeling for Level 3 automation. These factors collectively reduce the generalizability of our findings.

Our study is mainly based on a simple, decentralized approach to address this particular occurrence of control transitions. Smarter solutions are possible and should be deployed. At least two come to mind: (i) the vehicle may issue a ToR when still within the ODD so that the ToC can take place just at the end of the ODD (if this is known in advance by the vehicle), or (ii) traffic management in combination with vehicle-to-infrastructure communication (V2I) can provide an optimal schedule for AVs to performs their ToCs at specific times and positions. Our previous study [34] considered some of those aspects, especially with the objective to maintain automated driving as long as possible. Clearly, such management considerations could also apply in terms of traffic safety, but these were not specifically addressed in this work.

Finally, it is essential to emphasize that the use case presented in this study focuses solely on limited Level 3 operations up to 60 km/h, which was motivated by system approvals in accordance with UNECE regulations from 2021. The safety implications addressed in our analysis pertain to mandatory control transitions in accelerating dynamic traffic conditions, thus amplifying adverse ToC effects in such heterogeneous mixed-autonomy traffic conditions. The duration for which these potential safety implications will persist and realistically materialize in real traffic with higher AV shares remains challenging to predict from our perspective. A certification of Level 3 systems up to 130 km/h will supposedly come in a few years ahead, as UN Regulation 157 has been amended in 2023 [7]. However, it should be acknowledged that the requirements for automation to reliably ensure a safe control transition at such high speeds are significantly more challenging to achieve. Previous simulation studies have already indicated potential safety risks on high-speed motorways. Nevertheless, we believe that further research is imperative to comprehensively assess the safety implications of control transitions in large-scale, mixed-autonomy traffic.

Author Contributions: Conceptualization, R.A. and P.W.; methodology, R.A. and P.W.; software, R.A.; validation, R.A.; formal analysis, R.A.; investigation, R.A.; resources, R.A.; data curation, R.A.; writing—original draft preparation, R.A.; writing—review and editing, R.A. and P.W.; visualization, R.A.; supervision, P.W.; project administration, R.A.; funding acquisition, R.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: All of our descriptions should be sufficient enough to reproduce the data with the help of the referenced open-source simulator: SUMO, v1.19. We are not permitted to publish the actual scripts used for conducting our simulations due to DLR's policy.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AV	Automated Vehicle
ACC	Adaptive Cruise Control
ALKS	Automated Lane Keeping System
HGV	Heavy Goods Vehicle
LGV	Light Goods Vehicle
MDRAC	Modified Deceleration Rate to Avoid Crash
MRM	Minimum Risk Maneuver
MV	Manual Vehicle
ODD	Operational Design Domain
SSM	Surrogate Safety Measure
ToC	Transition of Control
ToR	Take-over Request
TTC	Time-to-Collision

References

1. Easy Tech. Easy Tech: Conditionally automated driving with the DRIVE PILOT. *Mag. Mobil. Soc.* **2021**, *12*. Available online: https://group.mercedes-benz.com/company/magazine/technology-innovation/easy-tech-drive-pilot.html (accessed on 23 October 2023).
2. Boeriu, H. BMWBLOG: BMW 7 Series Receives Approval Level 3 Automated Driving in Germany. *BMWBLOG* **2023**. Available online: https://www.bmwblog.com/2023/09/26/bmw-7-series-receives-approval-level-3-automated-driving-in-germany/ (accessed on 23 October 2023).
3. Lu, Z. Human Factors of Transitions in Automated Driving. Ph.D. Thesis, Delft University of Technology, Delft, The Netherlands, 2020. Available online: https://doi.org/10.4233/uuid:88dcb158-5fc3-4222-a402-4e484fa84414 (accessed on 1 November 2023).
4. Yu, D.; Nasir, M.; Pitts, B.J.; Shutko, J.; Bao, S.; Monk, C. L3 Vehicles are becoming a Reality: Important Human Factors Consideration for the Viability of Conditional Automation. *Proc. Hum. Factors Ergon. Soc. Annu. Meet.* **2023**, *67*, 1285–1288. [CrossRef]
5. On-Road Automated Driving (ORAD) Committee. In *Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles*; SAE International: Warrendale, PA, USA, 2018. [CrossRef]
6. United Nations Economic Commission for Europe (UNECE). Addendum 156—UN Regulation No. 157—Uniform Provisions Concerning the Approval of Vehicles with Regard to Automated Lane Keeping Systems. 2021. Available online:https://unece.org/sites/default/files/2021-03/R157e.pdf (accessed on 23 October 2023).
7. United Nations Economic Commission for Europe (UNECE). Addendum 156—UN Regulation No. 157—Amendment 4-Uniform Provisions Concerning the Approval of Vehicles with Regard to Automated Lane Keeping Systems. 2023. Available online: https://unece.org/transport/documents/2023/03/standards/un-regulation-157-amend4 (accessed on 28 November 2023).
8. Lücken, L.; Mintsis, E.; Porfyri, K.; Alms, R.; Flötteröd, Y.P.; Koutras, D. From Automated to Manual-Modeling Control Transitions with SUMO. In *Proceedings of the SUMO User Conference 2019*; Weber, M., Bieker-Walz, L., Hilbrich, R., Behrisch, M., Eds.; EPiC Series in Computing; EasyChair: Stockport, UK 2019; Volume 62, pp. 124–144. [CrossRef]
9. Maerivoet, S.; Akkermans, L.; Carlier, K.; Flötteröd, Y.P.; Lücken, L.; Alms, R.; Mintsis, E.; Koutras, D.; Wijbenga, A.; Vreeswijk, J.; et al. TransAID Deliverable 4.2-Preliminary Simulation and Assessment of Enhanced Traffic Management Measures. 2019. Available online: https://cordis.europa.eu/project/id/723390/results (accessed on 1 November 2023).
10. Merat, N.; Jamson, A.H.; Lai, F.C.; Daly, M.; Carsten, O.M. Transition to manual: Driver behaviour when resuming control from a highly automated vehicle. *Transp. Res. Part F Traffic Psychol. Behav.* **2014**, *27*, 274–282. . [CrossRef]
11. Lu, Z.; Happee, R.; Cabrall, C.; Kyriakidis, M.; de Winter, J. Human Factors of Transitions in Automated Driving: A General Framework and Literature Survey. *Transp. Res. Part F Traffic Psychol. Behav.* **2016**, *43*, 183–196. [CrossRef]
12. Gluck, A.; Deng, M.; Zhao, Y.; Menassa, C.; Li, D.; Brinkley, J.; Kamat, V. Exploring Driver Physiological Response During Level 3 Conditional Driving Automation. In Proceedings of the 2022 IEEE 3rd International Conference on Human-Machine Systems (ICHMS), Orlando, FL, USA, 17–19 November 2022; pp. 1–5. [CrossRef]
13. Jin, M.; Lu, G.; Chen, F.; Shi, X. How Driving Experience Affect Trust in Automation from Level 3 Automated Vehicles? An Experimental Analysis. In Proceedings of the 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), Rhodes, Greece, 20–23 September 2020; pp. 1–6. [CrossRef]
14. Vasile, L.; Dinkha, N.; Seitz, B.; Däsch, C.; Schramm, D. Comfort and Safety in Conditional Automated Driving in Dependence on Personal Driving Behavior. *IEEE Open J. Intell. Transp. Syst.* **2023**, *2023*, 772–784. [CrossRef]
15. Pipkorn, L.; Tivesten, E.; Flannagan, C.; Dozza, M. Driver Response to Take-Over Requests in Real Traffic. *IEEE Trans. Hum.-Mach. Syst.* **2023**, *53*, 823–833. [CrossRef]
16. Rangesh, A.; Deo, N.; Greer, R.; Gunaratne, P.; Trivedi, M.M. Autonomous Vehicles that Alert Humans to Take-Over Controls: Modeling with Real-World Data. In Proceedings of the 2021 IEEE International Intelligent Transportation Systems Conference (ITSC), Indianapolis, IN, USA, 19–22 September 2021; pp. 231–236. [CrossRef]

17. Neubauer, C.E.; Matthews, G.; De Los Santos, E.P. Fatigue and Secondary Media Impacts in the Automated Vehicle: A Multidimensional State Perspective. *Safety* **2023**, *9*, 11. [CrossRef]
18. Karakaya, B.; Bengler, K. Minimal Risk Maneuvers of Automated Vehicles: Effects of a Contact Analog Head-Up Display Supporting Driver Decisions and Actions in Transition Phases. *Safety* **2023**, *9*, 7. [CrossRef]
19. McDonald, A.D.; Alambeigi, H.; Engström, J.; Markkula, G.; Vogelpohl, T.; Dunne, J.; Yuma, N. Toward computational simulations of behavior during automated driving takeovers: A review of the empirical and modeling literatures. *Hum. Factors* **2019**, *61*, 642–688. [CrossRef]
20. Van Lint, J.; Calvert, S.C. A generic multi-level framework for microscopic traffic simulation—Theory and an example case in modelling driver distraction. *Transp. Res. Part B Methodol.* **2018**, *117*, 63–86. [CrossRef]
21. Calvert, S.; van Arem, B. A generic multi-level framework for microscopic traffic simulation with automated vehicles in mixed traffic. *Transp. Res. Part C Emerg. Technol.* **2020**, *110*, 291–311. [CrossRef]
22. Lopez, P.A.; Behrisch, M.; Bieker-Walz, L.; Erdmann, J.; Flötteröd, Y.P.; Hilbrich, R.; Lücken, L.; Rummel, J.; Wagner, P.; Wießner, E. Microscopic Traffic Simulation using SUMO. In Proceedings of the The 21st IEEE International Conference on Intelligent Transportation Systems, Maui, HI, USA, 4–7 November 2018; pp. 2575–2582.
23. Todosiev, E.P. *The Action Point Model of the Driver-VEHICLE System*; The Ohio State University: Columbus, OH, USA, 1963.
24. Xin, W.; Hourdos, J.; Michalopoulos, P.; Davis, G. The less-than-perfect driver: A model of collision-inclusive car-following behavior. *Transp. Res. Rec.* **2008**, *2088*, 126–137. [CrossRef]
25. Gardiner, C. *Stochastic Methods: A Handbook for the Natural and Social Sciences*, 4th ed.; Springer: Berlin/Heidelberg, Germany, 2009.
26. Treiber, M.; Kesting, A. *Traffic Flow Dynamics: Data, Models and Simulation*, 1st ed.; Springer: Berlin, Heidelberg, 2013. [CrossRef]
27. Favarò, F.; Eurich, S.; Nader, N. Autonomous vehicles': Trends, triggers, and regulatory limitations. *Accid. Anal. Prev.* **2018**, *110*, 136–148. [CrossRef]
28. Wang, S.; Li, Z. Exploring causes and effects of automated vehicle disengagement using statistical modeling and classification tree based on field test data. *Accid. Anal. Prev.* **2019**, *129*, 44–54. [CrossRef]
29. Boggs, A.M.; Arvin, R.; Khattak, A.J. Exploring the who, what, when, where, and why of automated vehicle disengagements. *Accid. Anal. Prev.* **2020**, *136*, 105406. [CrossRef]
30. Khattak, Z.H.; Fontaine, M.D.; Smith, B.L. Exploratory investigation of disengagements and crashes in autonomous vehicles under mixed traffic: An endogenous switching regime framework. *IEEE Trans. Intell. Transp. Syst.* **2020**, *2020*, 7485–7495. [CrossRef]
31. Agriesti, S.A.M.; Ponti, M.; Marchionni, G.; Gandini, P. Cooperative messages to enhance the performance of L3 vehicles approaching roadworks. *Eur. Transp. Res. Rev.* **2021**, *13*, 1–19. [CrossRef]
32. Kitajima, S.; Chouchane, H.; Antona-Makoshi, J.; Uchida, N.; Tajima, J. A Nationwide Impact Assessment of Automated Driving Systems on Traffic Safety Using Multiagent Traffic Simulations. *IEEE Open J. Intell. Transp. Syst.* **2022**, *3*, 302–312. [CrossRef]
33. Rösener, C.; Hennecke, F.; Sauerbier, J.; Zlocki, A.; Kemper, D.; Eckstein, L.; Oeser, M. A Traffic-based Method for Safety Impact Assessment of Road Vehicle Automation. In Proceedings of the Automated Vehicles Symposium, San Francisco, CA, USA, 9–12 July 2018; p. 14.
34. Alms, R.; Noulis, A.; Mintsis, E.; Lücken, L.; Wagner, P. Reinforcement Learning-Based Traffic Control: Mitigating the Adverse Impacts of Control Transitions. *IEEE Open J. Intell. Transp. Syst.* **2022**, *3*, 187–198. [CrossRef]
35. Wang, C.; Xie, Y.; Huang, H.; Liu, P. A review of surrogate safety measures and their applications in connected and automated vehicles safety modeling. *Accid. Anal. Prev.* **2021**, *157*, 106157. [CrossRef] [PubMed]
36. Das, T.; Samandar, M.S.; Autry, M.K.; Rouphail, N.M. Surrogate Safety Measures: Review and Assessment in Real-World Mixed Traditional and Autonomous Vehicle Platoons. *IEEE Access* **2023**, *11*, 32682–32696. [CrossRef]
37. Hayward, J.C. Near-miss determination through use of a scale of danger. *Highw. Res. Rec.* **1972**, *384*, 24-34.
38. Cooper, D.F.; Ferguson, N. Traffic studies at T-Junctions. 2. A conflict simulation record. *Traffic Eng. Control* **1976**, *17*, 306-309 .
39. Laureshyn, A.; De Ceunynck, T.; Karlsson, C.; Åse Svensson.; Daniels, S. In search of the severity dimension of traffic events: Extended Delta-V as a traffic conflict indicator. *Accid. Anal. Prev.* **2017**, *98*, 46–56. [CrossRef] [PubMed]
40. Ye, L.; Yamamoto, T. Evaluating the impact of connected and autonomous vehicles on traffic safety. *Phys. A: Stat. Mech. Its Appl.* **2019**, *526*, 121009. [CrossRef]
41. Ghanipoor Machiani, S.; Abbas, M. Safety surrogate histograms (SSH): A novel real-time safety assessment of dilemma zone related conflicts at signalized intersections. *Accid. Anal. Prev.* **2016**, *96*, 361–370. [CrossRef]
42. Vogel, K. A comparison of headway and time to collision as safety indicators. *Accid. Anal. Prev.* **2003**, *35*, 427–433. [CrossRef]
43. Kuang, Y.; Qu, X.; Weng, J.; Etemad-Shahidi, A. How Does the Driver's Perception Reaction Time Affect the Performances of Crash Surrogate Measures? *PLoS ONE* **2015**, *10*, e013861. [CrossRef]
44. Fazekas, A.; Hennecke, F.; Kalló, E.; Oeser, M. A Novel Surrogate Safety Indicator Based on Constant Initial Acceleration and Reaction Time Assumption. *J. Adv. Transp.* **2017**, *2017*, 8376572. [CrossRef]
45. Ozbay, K.; Yang, H.; Bartin, B.; Mudigonda, S. Derivation and Validation of New Simulation-Based Surrogate Safety Measure. *Transp. Res. Rec.* **2008**, *2083*, 105–113. [CrossRef]
46. Zheng, L.; Sayed, T.; Essa, M. Validating the bivariate extreme value modeling approach for road safety estimation with different traffic conflict indicators. *Accid. Anal. Prev.* **2019**, *123*, 314–323. [CrossRef] [PubMed]

47. Mintsis, E.; Koutras, D.; Porfyri, K.; Mitsakis, E.; Lücken, L.; Erdmann, J.; Flötteröd, Y.P.; Alms, R.; Rondinone, M.; Maerivoet, S.; et al. TransAID Deliverable 3.1-Modelling, Simulation and Assessment of Vehicle Automations and Automated Vehicles' Driver Behaviour in Mixed Traffic. 2019. Available online: https://cordis.europa.eu/project/id/723390/results (accessed on 1 November 2023).
48. Bolovinou, A.; Anagnostopoulou, C.; Roungas, V.; Amditis, A.; González, R.B.; Coello, L.T.; Álvarez González, D.; Kleinehagenbrock, M.; Neßler, J.; Rondinone, M.; et al. HI-DRIVE Deliverable D3.1/Use Cases Definition and Description. 2023. Available online: https://www.hi-drive.eu/app/uploads/2023/05/Hi-Drive-SP3-D3.1-Use-cases-definition-and-description-v1.1.pdf (accessed on 1 November 2023).
49. Sauvaget, J.L.; Dakil, M.; Griffon, T.; Anagnostopoulou, C.; Bolovinou, A.; Sintonen, H.; Metz, B. HI-DRIVE Deliverable D5.1/Descriptions of "Operations". 2023. Available online: https://www.hi-drive.eu/app/uploads/2023/05/Hi-Drive-SP5-D5.1-Description-of-Operations-v1.1.pdf (accessed on 1 November 2023).
50. Alms, R.; Flötteröd, Y.P.; Mintsis, E.; Maerivoet, S.; Correa, A. Traffic Management for Connected and Automated Vehicles on Urban Corridors-Distributing Take-Over Requests and Assigning Safe Spots. In Proceedings of the MFTS 2020 the 3rd Symposium on Management of Future Motorway and Urban Traffic Systems, Luxemburg, 6–7 July 2020.
51. Maggi, D.; Romano, R.; Carsten, O. Transitions Between Highly Automated and Longitudinally Assisted Driving: The Role of the Initiator in the Fight for Authority. *Hum. Factors* **2022**, *64*, 601–612. [CrossRef] [PubMed]
52. Eriksson, A.; Stanton, N.A. Takeover Time in Highly Automated Vehicles: Noncritical Transitions to and From Manual Control. *Hum. Factors* **2017**, *59*, 689–705. [CrossRef] [PubMed]
53. Papazikou, E.; Quddus, M.; Thomas, P.; Kidd, D. What came before the crash? An investigation through SHRP2 NDS data. *Saf. Sci.* **2019**, *119*, 150–161. [CrossRef]
54. Minderhoud, M.M.; Bovy, P.H. Extended time-to-collision measures for road traffic safety assessment. *Accid. Anal. Prev.* **2001**, *33*, 89–97. [CrossRef] [PubMed]
55. Martens, M.; Brouwer, R. Linking behavioral indicators to safety: What is safe and what is not? In Proceedings of the 3rd International Conference on Road Safety and Simulation (RSS 2011), Indianapolis, IN, USA, 14–16 September 2011.
56. Smirnov, N.V. Estimate of deviation between empirical distribution functions in two independent samples. *Bull. Mosc. Univ.* **1939**, *2*, 3–16.
57. Cucconi, O. Un nuovo test non parametrico per il confronto fra due gruppi di valori campionari. *G. Degli Econ. E Ann. Di Econ.* **1968**, *27*, 225–248.
58. Green, M. "How Long Does It Take to Stop?" Methodological Analysis of Driver Perception-Brake Times. *Transp. Hum. Factors* **2000**, *2*, 195–216. [CrossRef]
59. Milanés, V.; Shladover, S.E. Modeling cooperative and autonomous adaptive cruise control dynamic responses using experimental data. *Transp. Res. Part C Emerg. Technol.* **2014**, *48*, 285–300. [CrossRef]
60. Xiao, L.; Wang, M.; van Arem, B. Realistic Car-Following Models for Microscopic Simulation of Adaptive and Cooperative Adaptive Cruise Control Vehicles. *Transp. Res. Rec.* **2017**, *2623*, 1–9. [CrossRef]
61. Gunter, G.; Gloudemans, D.; Stern, R.E.; McQuade, S.; Bhadani, R.; Bunting, M.; Delle Monache, M.L.; Lysecky, R.; Seibold, B.; Sprinkle, J.; et al. Are Commercially Implemented Adaptive Cruise Control Systems String Stable? *IEEE Trans. Intell. Transp. Syst.* **2021**, *22*, 6992–7003. [CrossRef]
62. Ciuffo, B.; Mattas, K.; Makridis, M.; Albano, G.; Anesiadou, A.; He, Y.; Josvai, S.; Komnos, D.; Pataki, M.; Vass, S.; et al. Requiem on the positive effects of commercial adaptive cruise control on motorway traffic and recommendations for future automated driving systems. *Transp. Res. Part C Emerg. Technol.* **2021**, *130*, 103305. [CrossRef]
63. Makridis, M.; Mattas, K.; Anesiadou, A.; Ciuffo, B. OpenACC. An open database of car-following experiments to study the properties of commercial ACC systems. *Transp. Res. Part C Emerg. Technol.* **2021**, *125*, 103047. [CrossRef]
64. Mattas, K.; Albano, G.; Donà, R.; Galassi, M.C.; Suarez-Bertoa, R.; Vass, S.; Ciuffo, B. Driver models for the definition of safety requirements of automated vehicles in international regulations. Application to motorway driving conditions. *Accid. Anal. Prev.* **2022**, *174*, 106743. [CrossRef]

3.4 Paper III - Traffic Capacity Constraints from Level 3 Control Transitions

Originally published as:

Alms R and Wagner P (2025) Traffic capacity constraints from level 3 control transitions. *Front. Future Transp.* 6:1600739. doi: 10.3389/ffutr.2025.1600739

Scope and Contribution of Paper III

As Papers I and II already touch on capacity-related effects without precise quantification, the third paper quantifies capacity limits across ToC scenarios with increasing AV shares on a two-lane motorway to address RQ 3. It further examines the mechanisms behind these capacity reductions using simplified numerical experiments that isolate the impact of the ToC preparation phase. First, a baseline SUMO scenario is calibrated to real-world detector data, then four scenarios are contrasted (no ToCs; unmanaged ToCs; coordinated ToC management; unmanaged ToCs "rightmost95"[9]). Subsequently, capacity estimators based on platoon experiments are compared against estimators using fixed headway assumptions for Level 3 AVs. The principal mechanism identified in these numerical experiments is an increase in peak time headways during quasi-synchronous ToCs; estimator variants that embed these peaks reproduce the overall trend better than fixed-headway assumptions.

Results from the simulation study show substantial capacity losses that increase with Level 3 share when ToCs are unmanaged. Coordinated TM upstream of the No-AD zone mitigates these losses, whereas the "rightmost95" case exacerbates them. Paper III consolidates earlier findings, showing that traffic efficiency deteriorates alongside ToC-induced capacity reductions and that the tested countermeasures do not restore baseline capacity, answering RQ 3 by quantifying the constraint and its limited recoverability [10]. It concludes by contrasting these results with optimistic claims in the literature regarding capacity gains from Level 3 or 4 AVs. Section 4.3 discusses these outcomes further with additional analyses of time headway and lane-change data.

[9]This scenario emulates the concept of the latest Level 3 version of Mercedes-Benz's Drive Pilot, operating up to 95 km/h in the rightmost lane without overtaking in car-following.

[10]Additionally, Section 4.2 reports SSM results from this simulation study, which were not included in the published paper, indicating detrimental safety impacts in mid-range AV shares.

Minor correction note:

The text in the second paragraph of section 3.1 of this paper states a time range from 2015 – 2022 regarding Figure 4b. The panels in this referenced subfigure show more recent plots from 2017 – 2024 though.

frontiers | Frontiers in Future Transportation

TYPE Original Research
PUBLISHED 07 July 2025
DOI 10.3389/ffutr.2025.1600739

OPEN ACCESS

EDITED BY
Aleksandar Stevanovic,
University of Pittsburgh, United States

REVIEWED BY
Haifei Yang,
Hohai University, China
Zeynel Baran Yıldırım,
Adana Science and Technology University, Türkiye
Changshuai Wang,
Southeast University, China
Jana Sarran,
University of Guyana, Guyana

*CORRESPONDENCE
Robert Alms,
Robert.Alms@dlr.de

RECEIVED 26 March 2025
ACCEPTED 16 June 2025
PUBLISHED 07 July 2025

CITATION
Alms R and Wagner P (2025) Traffic capacity constraints from level 3 control transitions. *Front. Future Transp.* 6:1600739. doi: 10.3389/ffutr.2025.1600739

Traffic capacity constraints from level 3 control transitions

Robert Alms [1]* and Peter Wagner [1,2]

[1]Institute of Transportation Systems, German Aerospace Center (DLR), Berlin, Germany, [2]Institute of Land and Sea Transport Systems, TU Berlin, Berlin, Germany

With the increasing integration of conditionally automated Level 3 systems into real-world traffic, concerns about their impact on traffic efficiency and capacity have emerged. When such systems reach their operational limits, mandatory control transitions could disrupt traffic flow and reduce overall capacity. This study employs large-scale simulations and numerical experiments to analyze these effects and quantify potential capacity constraints. The results of the two-lane highway scenario show an experimental capacity reduction of up to 2000 veh/h in an almost fully automated but unmanaged traffic mix, corresponding to a loss of about 60%. Control transition-related effects become increasingly pronounced at a Level 3 penetration rate between 10% and 20%. Estimated capacity reductions suggest that the maxima in time headway increments during the transition phase contribute most to these effects.

KEYWORDS

automated vehicles (AVs), level 3 automation, mixed-autonomy traffic, traffic capacity, transition of control (ToC)

1 Introduction

As manufacturers begin to introduce Level 3 automated driving systems to the market, the potential impact of such systems on overall traffic flow and capacity needs to be investigated. A key challenge arises from the fact that Level 3 systems require human drivers to take over control when reaching system limits, leading to so-called transitions of control (ToC), which may disrupt traffic flow and reduce road capacity. Despite regulatory advancements concerning Level 3 systems (R157 by UNECE (2023)), the macroscopic impact of such procedural ToC effects on traffic conditions remains insufficiently explored. This raises the general question of how Level 3 control transitions in conditionally automated vehicles (AVs) affect traffic capacity and, more specifically, what characteristics of procedural ToC-induced time headway increments in vehicle strings contribute to this effect. To investigate this, we conduct a large-scale simulation-based analysis and complement it with simplified numerical experiments to estimate macroscopic capacity impacts. Our study also explores the underlying mechanisms of the transition phase in greater detail. Existing research on potential capacity gains from AVs, as exemplified by Friedrich (2016) and Park et al. (2021), has primarily focused on higher automation levels (4–5) under optimistic assumptions of short time headways, e.g., $\tau_{AV} = 0.5$ s, in contrast to observed headways in manually driven vehicles of at least 1 s in freeway traffic, depending on vehicle speed, as shown by Wagner (2012). Our previous work in Alms et al. (2022) and Alms and Wagner (2024) touched on ToC-related capacity effects but lacked a comprehensive quantification of resulting capacity losses. This study addresses these gaps by (i) adopting a macroscopic perspective using realistic, R157-compliant time headways of $\tau_{AV} = 1.6$ s, and (ii) introducing an exploratory estimation approach that explicitly accounts for Level 3 disengagements in road capacity assessment.

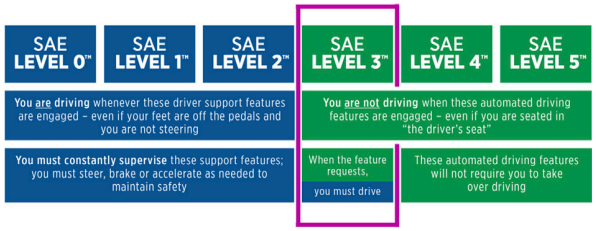

FIGURE 1
Excerpt from Shuttleworth (2019)'s illustration of the SAE Levels of Driving Automation (SAE International, 2021), with Level 3 highlighted by a purple frame.

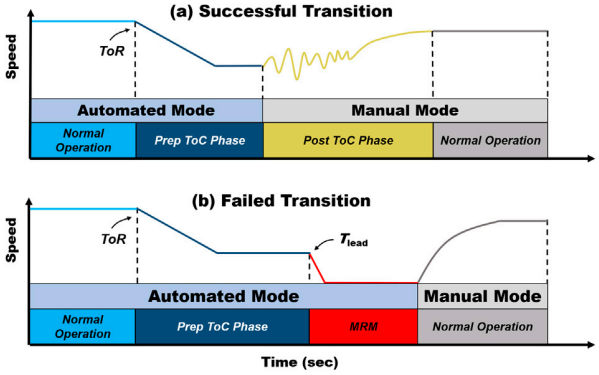

FIGURE 2
ToC model operation modes illustrated in a representative speed-time diagram for: **(a)** successful transition and **(b)** failed transition.

The rest of the paper is organized as follows: Section 2 introduces the conceptual aspects of ToCs in Level 3 automated systems. In Section 3, we present a highway scenario calibration based on real-world detector data. Section 4 details our methodology for investigating ToC-related capacity effects in a simulation study, while Section 5 presents and discusses our results, comparing simulated and estimated capacity reductions. Lastly, Section 6 offers our perspective on the interpretation and limitations of this study.

2 Transitions of control in level 3 automated driving

The six levels of driving automation, defined by SAE International (2021), not only classify automated driving functions and capabilities but also specify the human driver's role in terms of engagement and responsibility, as illustrated in Figure 1. Conditional automated driving (Level 3, highlighted with a purple frame in Figure 1) represents a fundamental shift toward automated vehicle operation within defined Operational Design Domains (ODD), specified in British Standards Institution (2020), allowing human drivers to disengage from the primary driving task. However, if the Level 3 system requires the driver to resume control, a takeover request (ToR) is issued, initiating a critical transfer of authority: these procedures are referred to as transitions of control (ToC, plural: ToCs). Detailed insights into various aspects of ToCs are available through a comprehensive literature review on takeovers in automated driving (McDonald et al., 2019). Further studies examine the intricacies of modeling human factors, such as situational awareness and task demand (Van Lint and Calvert, 2018; Calvert and van Arem, 2020), or reduced driver performance (Wang et al., 2025b), during ToCs.

The current regulations R157 from UNECE (2023) specify technical requirements for the certification of Level 3 Automated Lane Keeping Systems (ALKS) and set the time range T_{lead} to 10 s before a failed transition escalates to a minimum risk maneuver (MRM), which is critical for the process of control transitions. Within the context of the EC project (TransAID, 2021; Lücken et al., 2019; Mintsis et al., 2019) introduced a novel ToC model, which is fully parametrizable to align with these later-established UNECE specifications and for which a detailed description of the model's implementation is provided. The operationalization of ToCs is further specified in the ongoing EC project Hi-Drive (Bolovinou et al., 2023; Sauvaget et al., 2023) and demonstrated in Schulte-Tigges et al. (2023).

Figure 2 illustrates the basic mechanisms of the ToC model for successful and failed control transitions implemented in the microscopic traffic simulation SUMO (Alvarez Lopez et al., 2018). After a ToR, the AV enters a preparatory phase characterized by headway enlargement and disabled lane changing. Automated driving continues for the limited lead time, after which either the driver resumes control in time (successful transition), or, if not, the AV initiates an MRM (failed transition). For failed transitions, the AV initiates a phase of constant deceleration and may come to a full stop if the human driver does not respond. Although such events are rare, they can have a high impact and are the subject of extensive safety investigations based on disengagement reports (e.g., (Ward, 2024; Kohanpour et al., 2025)). However, this aspect is not the focus of the present work. In the case of a successful transition, the driver state model accounts for a phase of reduced human driving performance, with recent studies gaining further insights into both post-ToC durations (Wang et al., 2025a) and potential negative impacts on traffic stability (Wang et al., 2024).

In Maerivoet et al. (2019) and Lücken et al. (2019) principal transition phase effects of consecutive, quasi-synchronous ToCs in a platoon of Level 3 automated vehicles were previously demonstrated. Figure 3a, which depicts speed and time headways for a string of five AVs disengaging at the same location, illustrates this effect in a simplified simulation experiment with identical vehicle parametrization. The increased time headways, and consequently the cumulative speed reduction, are caused by the preparatory headway increment of the vehicle automation to facilitate a safe takeover (*cf.* Figure 2, *Prep ToC Phase*). Figure 3b extends this analysis by showing acceleration profiles for a larger platoon of up to 32 vehicles—the maximum size at which the last AV still manages to prevent a complete stop—using SUMO's ACC model for AVs, based on Xiao et al. (2017). The main observed effects in the vehicle decelerations include:

- With a moderate default deceleration of 1 m/s^2 during the 10 s transition phase specified by R157, maintaining safe gaps in

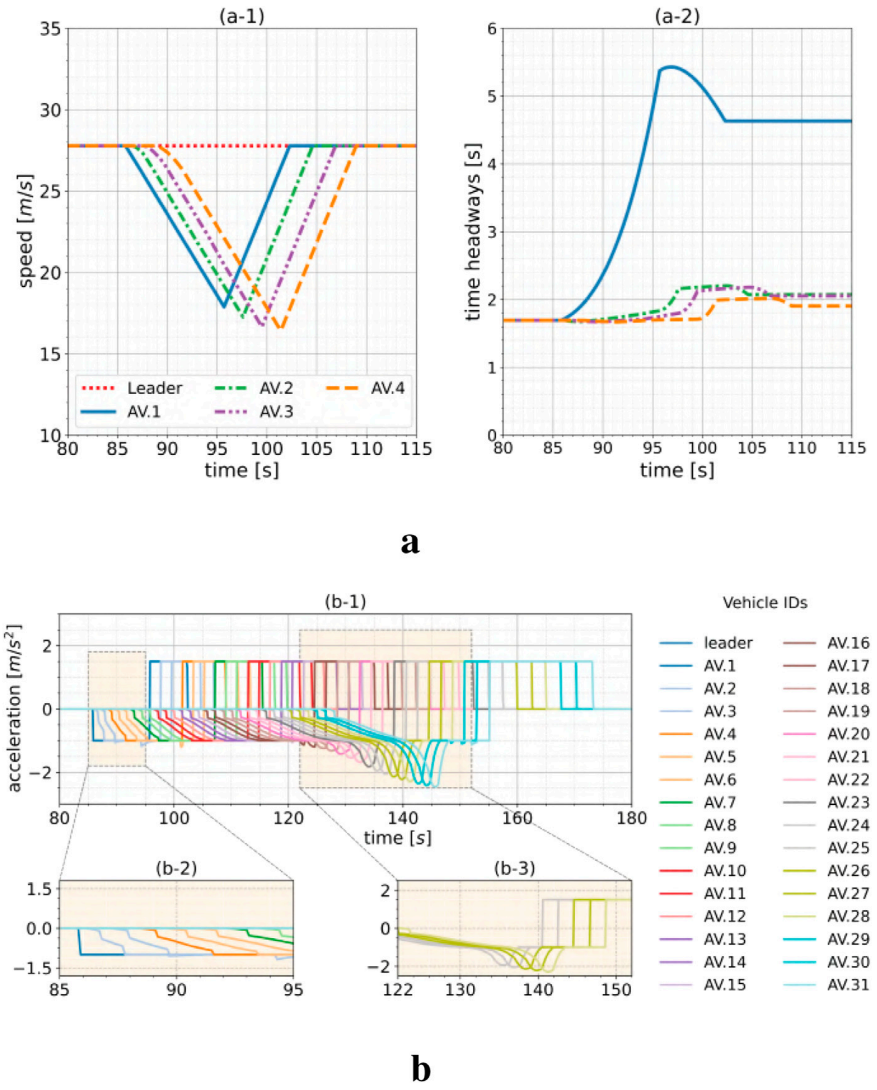

FIGURE 3
Platoon simulations with consecutive ToCs performed at a fixed location. **(a)** shows speed and time headway profiles for a short platoon with five AVs, illustrating the headway increment effect. **(b)** depicts acceleration profiles for a platoon of 32 AVs, showing an escalating headway increment effect with SUMO's ACC model.

AV platoons is not feasible without initiating deceleration earlier. Panel (b-2) in Figure 3b illustrates that, starting with the first vehicle behind AV.1 (dark blue line), SUMO's gap controller begins to decelerate even before the respective vehicles receive ToRs to initiate their ToC.

- Starting with AV.20, the following vehicles must decelerate more aggressively than their target deceleration of 1 m/s^2. Panel (b-3) in Figure 3b highlights these deceleration overshoots for AV.24–28. These overshoots are specific to the ACC model, while similar experiments employing SUMO's default model do not exhibit this behavior. However, that model compensates by initiating deceleration even earlier than the ACC model. The principal accumulation effect of consecutive ToCs remains present in both cases.

These numerical experiments are highly simplified due to identical vehicle parametrizations, yet they effectively illustrate the isolated ToC effects discussed. Given the cumulative deceleration patterns observed, we expect that ToC-induced disturbances may lead to noticeable reductions in traffic capacity. To examine whether these effects also manifest under more realistic traffic flow conditions, we calibrate a SUMO simulation scenario to detector data in Section 3.

3 Calibrating SUMO for a highway traffic scenario

To analyze the impact of ToCs on traffic capacity, we use real-world detector data from a German highway west of Berlin as a

reference for SUMO calibration. The following sections detail the dataset and simulation setup.

3.1 AVUS detector data

Figure 4a shows a section of the Bundesautobahn A115, referred to as AVUS, which was occasionally used as a motor racing track in the past and is a highly frequented highway with up to 80.000 vehicles per day. We hypothesize a potential ODD zone for Level 3 automated driving in the inbound segment of the road (*cf.* Figure 4a, panel (a)), which is a two-lane highway with speed limits of 80 km/h starting at an interchange section and increasing to 100 km/h up to the next traffic exit, that is located nearly 5 km downstream. The inbound traffic data on the two-lane section come from a detector at an underpass (*cf.* Figure 4a, panel (c)). After this point, the road has a slight slope for a few hundred meters, but the exact gradient could not be verified. A ramp merges onto the main edge about 600 m downstream, from where the speed limit increases to 100 km/h.

Figure 4b displays speed—flow relations for several years of the AVUS between 2015 and 2022, as scatterplots based on data from Digitale Plattform Stadtverkehr Berlin (2024). These data are originally tagged as hourly flows with corresponding average speeds per hour, but we suspect that this is not accurate. While the number of vehicles is accumulated over a full hour, the high variations in speeds at lower flow rates suggest that these data points from the detector database might actually represent speed averages over intervals of 1 minute or less. We were unable to verify this suspicion directly with the publisher of the data, but we argue that the actual speed value recorded in the database is likely the last entry of a full hour — possibly for efficiency and memory-saving reasons in data processing — rather than the average speed over the entire hour. This ultimately results in a notably wider distribution of speed values at lower flow rates than expected for true hourly data. For reference, we also added the model fit developed by Van Aerde (1995) to each plot.

Table 1 lists the yearly maximum flows q, the 95th and 99th percentiles as suggested by Brilon and Geistefeldt (2010), and the deterministic capacity derived from the van Aerde model, as well as the corresponding shares for heavy good vehicles (HGVs) extracted from the raw detector data between 2016 and 2024. Additionally, data provided by the BASt (2025) from a detector downstream of the AVUS at "Eichkamp" are also included in the table for comparison (*cf.* Figure 4a, panel (b)). Note that 2024 shows an oddly high HGV share, which we consider very unlikely and attribute to recent technical changes in sensor-based detection and data processing by the provider. The report from BASt (2021) stated a nationwide HGV share of 18.1% in 2021 on Germany's highways.

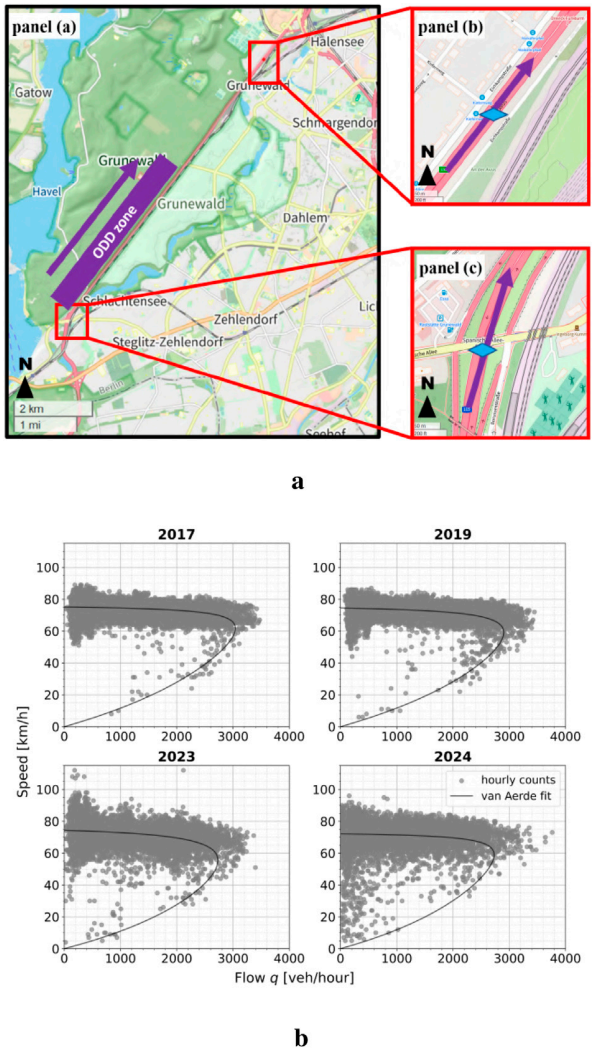

FIGURE 4
Overview of the AVUS scenario. **(a)** OpenStreetMap view of the AVUS highway with a potential 5 km ODD zone on the two-lane inbound edge highlighted in purple. Zoom displays show traffic detector locations, marked with blue in subpanels **(b)** and **(c)**. Purple arrows indicate the inbound traffic direction. (b) Speed–flow scatterplots of yearly AVUS detector data for inbound traffic from detector "TE002".

3.2 Simulation setup for calibration

To investigate the impacts of ToCs in mixed-autonomy traffic, we compose a traffic mix of four different vehicle types: automated passenger vehicles (AVs), manual passenger vehicles (MVs), light goods vehicles (LGVs), and heavy goods vehicles (HGVs). The most relevant parameters for a heterogeneous traffic behavior in this AVUS highway scenario are visualized in Figure 5. Instead of utilizing SUMO's default parameters, vehicle type specific distributions were deployed. Table 2 presents the full parametrization scheme for all vehicle types.

In principle, SUMO's vehicle insertion capacity exceeds that of comparable real-world traffic scenarios. Therefore, we aim to calibrate the simulation primarily to match the maximum flow q in relation to the real-word AVUS data. Besides the general vehicle parametrization, two insertion properties in SUMO heavily effect the overall capacity of a simulation, i.e., the vehicle speed at insertion `departSpeed` and lane choice at insertion `departLane`. We kept these parameters unchanged for all simulations in the paper. The most important capacity related SUMO options are defined as follows:

- `departSpeed = max`
- `departLane = random (AV,MV,LGV)`

TABLE 1 Yearly flow metrics and HGV shares for AVUS inbound traffic.

Year	2016	2017	2018	2019	2020	2021	2022	2023	2024
Max q	3,477	3,472	3,497	3,447	3,528	3,226	3,284	3,396	3,763
99%ile	3,212	3,155	3,116	3,142	3,104	2,865	2,837	2,943	3,018
95%ile	2,818	2,751	2,708	2,766	2,677	2,434	2,530	2,514	2,541
van Aerde c_F	3,070	3,011	3,013	2,915	2,971	2,687	2,622	2,733	2,766
raw HGV (%)	5.94	5.74	5.92	5.55	5.60	4.91	6.37	7.72	*28.46
BASt HGV (%)	6.51	7.21	7.25	6.81	7.24	7.56	7.05	—	—

*Outlier value; see main text for discussion

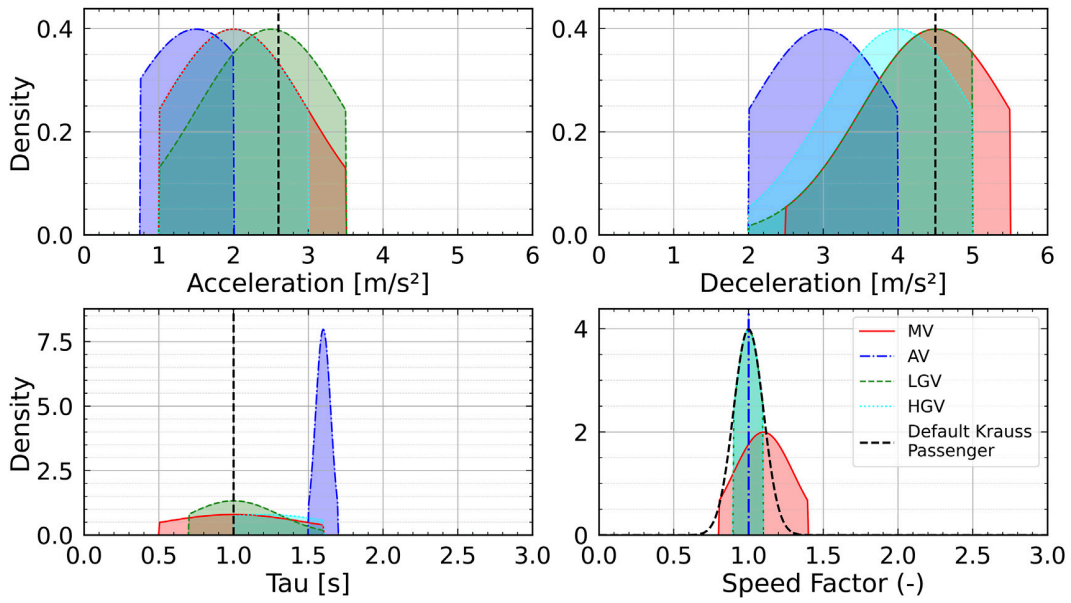

FIGURE 5
Distributions of the main parameters for different vehicle types. SUMO's default is indicated by the black dashed line.

- `departLane = right (HGV)`
- `extrapolate-departpos = true`
- `step-length = 0.1 s`

3.3 Calibration including ramp flow

In the first step, we ran simulations with increasing demand using MVs only. To better capture the full spectrum of the fundamental diagram in SUMO, we introduced additional vehicle flow on the incoming ramp. This creates a merging scenario, leading to traffic breakdown upstream of the main edge's detector position. Figure 6 shows the speed–flow relations as scatterplots for (i) real-world detector data from 2024, (ii) SUMO's default parametrization, and (iii) the aggregated main edge data from the final calibration. The graphic also color-codes the demand intensities from the on-ramp and highlights the maximum q of the AVUS detector data 2024. For reference of the expected average speeds defined by the german Highway Capacity Manual—referred to as HBS—in (FGSV, 2015, Part A, Figure A3-10) for a two-lane highway (slope ≤ 2, speed limit 80 km/h), a black solid line was added.

The key findings derived from Figure 6 are:

1. Comparing the dark gray SUMO default data with the light gray detector data, we identify how far SUMO's default exceeds the actual maximum flow (about 700 vehicles surplus).
2. The calibrated main edge's flow (blue-colored points) is notably lower than SUMO's default. Maximum flows (dark blue points) are much closer to the real-data (about 130 vehicles difference) compared to SUMO's default (dark gray).
3. The overall speed–flow relation of the calibrated main edge (blue) is slightly tilted toward higher speeds compared to SUMO's default (dark gray), and the speed gradient more in line of the HBS expectation (black line).
4. The calibrated main edge's traffic breakdown on the congested side of the fundamental diagram (indicated by darker-colored blue points) is much less pronounced than what is to be expected from real-world data (see light gray scatter points).

TABLE 2 Vehicle type definitions. "—" indicates not defined.

Parameter/Attribute	MV	LGV	HGV	AV
Car-following model	Krauss	Krauss	Krauss	ACC
sigma	$\mathcal{N}(0.2, 0.50)$, [0, 1]	$\mathcal{N}(0.1, 0.20)$, [0.0, 1.0]	$\mathcal{N}(0.1, 0.20)$, [0.0, 1.0]	—
tau	$\mathcal{N}(1.0, 0.50)$, [0.5, 1.6]	$\mathcal{N}(1.0, 0.30)$, [0.7, 1.6]	$\mathcal{N}(1.2, 0.50)$, [1.0, 1.6]	$\mathcal{N}(1.6, 0.05)$, [1.5, 1.7]
decel	$\mathcal{N}(4.5, 1.00)$, [2.5, 5.5]	$\mathcal{N}(4.5, 1.00)$, [2.0, 5.0]	$\mathcal{N}(4.0, 1.00)$, [2.0, 5.0]	$\mathcal{N}(3.0, 1.00)$, [2.0, 4.0]
accel	$\mathcal{N}(2.0, 1.00)$, [1.0, 3.5]	$\mathcal{N}(2.5, 1.00)$, [1.0, 3.5]	$\mathcal{N}(2.0, 1.00)$, [1.0, 3.0]	$\mathcal{N}(1.5, 1.00)$, [0.75, 2.0]
speedFactor	$\mathcal{N}(1.1, 0.20)$, [0.8, 1.4]	$\mathcal{N}(1.0, 0.10)$, [0.9, 1.1]	$\mathcal{N}(1.0, 0.10)$, [0.9, 1.1]	1.0
lcAssertive	$\mathcal{N}(1.3, 0.40)$, [0.9, 1.7]	$\mathcal{N}(1.1, 0.05)$, [1.0, 1.1]	$\mathcal{N}(1.0, 0.05)$, [0.9, 1.1]	$\mathcal{N}(0.7, 0.10)$, [0.6, 0.8]
vClass	Passenger	Delivery	Truck	Passenger
length [m]	5.0	8.0	15.0	5.0
width [m]	1.8	2.0	2.4	1.8
actionStepLength [s]	0.1	0.1	0.1	0.1
maxSpeed [m/s]	55.56	27.78	25.0	55.56
speedDev	0	0	0	0.01
TOC model—moderate parametrization scheme				
TOC device	—	—	—	true
manualType	—	—	—	MV
automatedType	-	—	—	AV
responseTime	—	—	—	$\mathcal{N}(7.0, 2.50)$, [2, 60]
initialAwareness	—	—	—	$\mathcal{N}(0.5, 0.30)$, [0.1, 1.0]
recoveryRate	—	—	—	$\mathcal{N}(0.2, 0.10)$, [0.01, 0.5]
mrmDecel	—	—	—	3.0
ogNewSpaceHeadway	—	—	—	10.0
ogNewTimeHeadway	—	—	—	5.0
ogChangeRate	—	—	—	1.0
ogMaxDecel	—	—	—	1.0

The phenomenon described in point 4 is, in part, a limitation of SUMO's current modeling of cooperative lane-changing behavior between neighboring lanes under traffic breakdown conditions. Correspondingly, Figure 7 compares the lane-specific calibrated flows in SUMO with real AVUS data from 2024. We clearly identify the disparate speed levels between the lanes in SUMO (bottom panel), whereas the real-world data (top panel) indicate similar speed–flow relations on both lanes. Rummel (2017) indirectly revealed this issue in his investigation but was unable to unequivocally identify the lane-specific breakdowns as the underlying cause of SUMO's oversaturation compared to the HBS predictions, nor did the report by Geistefeldt et al. (2017), which ultimately disregarded SUMO in its analysis for this very reason. While this limitation prevents a full replication of the real-world dynamics, we proceed with the calibration of the scenario as a basis for our analysis and will address this shortcoming in our future work.

3.4 Refining calibration by incorporating HGV share

In a second step, based on the parametrization scheme plausibilised for MVs in the ramp scenario (*cf.* Table 2), we conducted simulations with different shares of HGVs, LGVs, and MVs, but without any ramp flow. As a result, we can no longer reproduce the entire fundamental diagram for this highway scenario, since SUMO's flow does not naturally lead to a traffic breakdown as observed in real-world highway traffic. The reason we need to disregard the unstable part of the fundamental diagram at this point is technical: SUMO does not maintain precise LGV/HGV shares for vehicle insertions when approaching maximum flow. Instead, the share of LGVs and HGVs declines to zero until SUMO can only insert MVs when the traffic breakdown at capacity is expected. This behavior stems partly from the parametrization of LGVs and HGVs, such as their larger vehicle lengths and time headways.

Figure 8 shows the speed–flow relations for the main edge with LGV/HGV shares of 0%, 5%, 10%, and 15% (LGV/HGV distributed as 2/3 vs. 1/3), compared to AVUS detector data from 2018. The key findings are as follows:

1. The maximum flow with a 0% HGV share (purple-colored scatter points) is almost the same as in the ramp case (*cf.* Figure 6, purple markers), with a difference of about 50 veh/h.
2. The AVUS detector data from 2018 have an HGV share of 6%, with a maximum flow of 3497 veh/h. The speed variation in the detector data, particularly for lower demands, is very high, which we attribute to factors discussed in Section 3.1.
3. The speed variations in all simulation results are relatively large. This is expected, as we deliberately plotted only the average speeds of the last 1-min interval of a full hour, which we suspect is also the case for the real detector data. This illustrates a plausible speed distribution from the calibrated simulations compared to the detector data.
4. The color-coded flows at capacity decrease notably with increasing HGV shares (decline by 260 to 410 veh/h).

Considering the relatively low deterministic capacities based on the AVUS detector data stated in Table 1 compared to the expected capacities from the HBS (range from 3,600 to 3900 veh/h) for this highway type, we assess our SUMO calibration in terms of capacity as follows:

- The maximum flow in the SUMO ramp scenario without HGV share consideration is 3907 veh/h. The 95th and 99th percentile flows are 3853 veh/h and 3882 veh/h, respectively. The van Aerde model estimates a maximum flow of 3665 veh/h based on SUMO data. These values are significantly higher than the detector data but do not account for HGV shares in SUMO.
- With HGV share consideration the maximum flows on the stable arm of the fundamental diagram decrease notably between 6.7 to 10.5% as illustrated in Figure 8.

Under the assumption that those percentages under HGV consideration scale down proportionally in SUMO with the capacity numbers stated above, we obtain the following deterministic capacity ranges for the calibrated parametrization scheme:

- Max flow: 3497 – 3647 veh/h
- 95th percentile: 3449 – 3597 veh/h
- 99th percentile: 3475 – 3624 veh/h
- van Aerde model: 3281 – 3421 veh/h

Even though these capacities are still about 200 – 300 veh/h larger than the detector numbers in Table 1, we consider this an adequate calibration, particularly compared to SUMO's default, since the real-world detector flow data are overall lower compared to the HBS range, which we identified to be between 3,600 and 3900 veh/h. Other local factors, such as road curvature, slope, shoulder lane width, underpass length, or surface conditions, which might impact the local capacity and could explain the rather low detector-based flows, are unknown to us. While more detailed microscopic calibration using high-resolution trajectory data (e.g., as in Schrader (2024) or Liu et al. (2024)) would be desirable, such data were not available for this study.

4 Methodology to quantify capacity effects of ToCs

To determine ToC-related capacity impacts, we conduct a simulation study with an increasing AV penetration rate and measure the corresponding maximum flows q. Additionally, we estimate the anticipated ToC-induced capacity reduction and later compare these estimates with the measured results from the simulation study.

4.1 Simulation experiment

```
Initialize:
  low ← 0, high ← max demand
  max_valid ← −1
  seeds ← 12
  threshold ← seeds/2 = 6
while low ≤ high do
    mid ← (low + high)/2
    Run simulation at demand mid for each seed in seeds
    Count valid and invalid results
     if invalid results ≤ threshold then
        max_valid ← mid
        Increase low
        Record maximum flow at detector for valid results
     else
        Decrease high
     end if
 end while
 Save max valid demand and corresponding maximum flow
```

Algorithm 1. Binary search for maximum flow.

For the simulation study, we define a wide range of traffic shares based on the vehicle types outlined in Table 2. The traffic compositions feature increasing AV shares (AV00–AV85) in 10% increments, with AV increases and MV decreases of equal size, and a constant LGV/HGV share of 15% (split 2/3 LGV, 1/3 HGV). Considering a hypothetical ODD zone on the AVUS inbound highway, as illustrated in Figure 4a, AVs are assumed to be capable of Level 3 automated driving at speeds of up to 100 km/h until reaching the end of the ODD zone. Vehicles enter the network in their respective driving mode at random on one of the two lanes, except for HGVs, which are only inserted on the right lane. They continue their trip until reaching the end of the AVUS, near the detector position highlighted in Figure 4a, panel (b). Four distinct scenarios are examined, differing in how ToCs are facilitated:

1. No ToCs: Simulations without any ToCs.

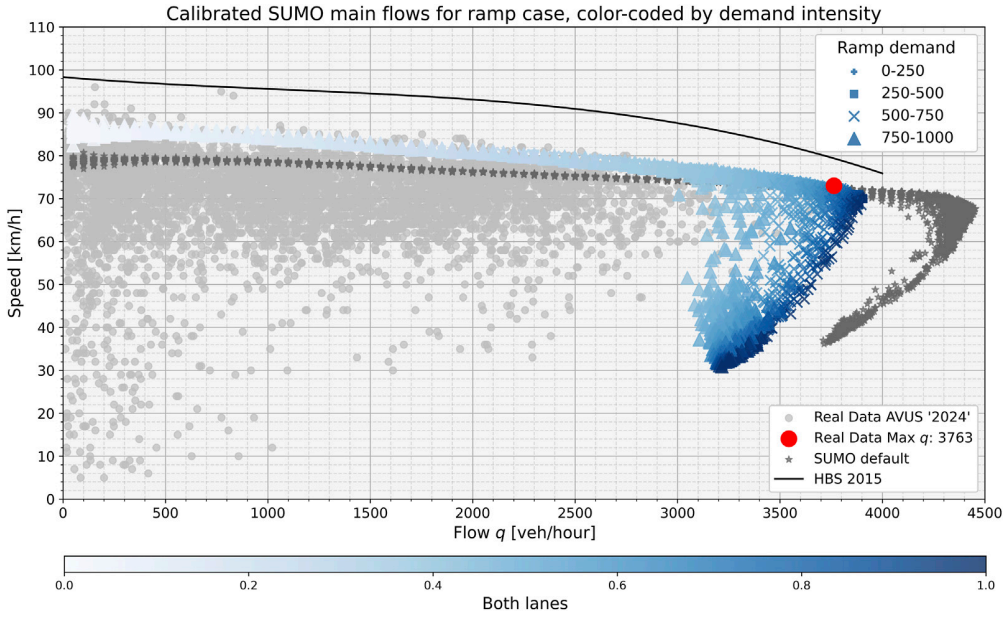

FIGURE 6
Speed–flow relations for the AVUS scenario, including ramp flow, comparing calibrated flows based on the parametrization scheme from Table 2 versus AVUS detector data from 2024, and SUMO's default. The main edge's demand intensities are color-coded in blue. Ramp demand levels are coded by marker size and style.

2. Unmanaged: Simulations with unmanaged ToCs at the end of the ODD zone.
3. Managed: Simulations with ToCs managed by a ToC-dispatch algorithm over the full length of the ODD zone.
4. Unmanaged "rightmost95": Simulations with unmanaged ToCs at the end of the ODD zone, emulating the concept of the latest approved manufacturer system by Mercedes-Benz Group (2024), operating up to 95 km/h on the rightmost lane, without overtaking.

For cases 2–4, we additionally run simulations with τ-distributions for MVs around $\tau_{MV} = 0.8$ s and 1.2 s. In case 3, we deploy the heuristic algorithm developed by Lücken et al. (2019). The control algorithm basically emulates a V2X-based traffic management scheme by dispatching ToRs to AVs in a coordinated manner to mitigate the accumulation effect of consecutive ToCs. For case 4, AVs are only inserted into the simulation on the rightmost lane, overtaking is disabled, and their speed is limited to 95 km/h.

To measure the capacity per AV share as precisely as possible, we run the AVUS scenario with 12 seeds per traffic mix, deploying a binary search as illustrated in Algorithm 1. To ensure we obtain the correct maximum flow, the results of each run must be checked against the actual traffic share versus the expected share due to SUMO's insertion mechanism, as described in Section 3.4. The binary search continues increasing the demand as long as valid traffic shares are observed, until the maximum flow per simulation run is reached.

Simulations run with a 30 min warm-up phase to populate the scenario and then record data for a full hour of simulated time (SUMO version 1.22 from Alvarez Lopez et al. (2025)). A detector near the end of the ODD zone records speed and flow to identify potential traffic breakdowns and measures the maximum flow. Figure 9 exemplarily shows spatiotemporal heatmaps of the ODD zone for speed and flow. Figures 9A,B, result from the same demand level and AV share—only the seed values, which determine the randomization process of vehicle insertions, differ.

4.2 Estimating capacity reduction

Considering the `minGap` in SUMO as $g_{\min}$, individual vehicle lengths l_i, type-specific τ_i, and varying vehicle shares p_i, the theoretical lane capacity C at speed v is given by:

$$C = \frac{v}{\sum_i p_i \cdot (l_i + g_{\min}) + v \sum_i p_i \cdot \tau_i} \tag{1}$$

For the SUMO default `minGap` of 2.5 m, a speed $v = 100\,km/h$, and the respective vehicle lengths and τ-values from Table 2, we compute lane capacities across all traffic mixes. Assuming a constant time headway τ_i for each vehicle type, without considering ToCs, the resulting capacities per mix are shown in Figure 10a. The results demonstrate that as the τ-values for MVs increase ($\tau_{\mathrm{MV}} = 0.8$ s to $\tau_{\mathrm{MV}} = 1.2$ s), while keeping fixed values for AVs ($\tau_{\mathrm{AV}} = 1.6$ s), LGVs ($\tau_{\mathrm{LGV}} = 1.0$ s), and HGVs ($\tau_{\mathrm{HGV}} = 1.2$ s), the decline in maximum capacity across traffic mixes becomes less pronounced. If MVs had the same τ-value as AVs—in this case, $\tau_{\mathrm{AV}} = 1.6$ s — the lane capacities would remain stable, regardless of the increasing AV share.

To account for ToC effects in such estimations, we repeat the simplified numerical experiment with the 32-vehicle platoon

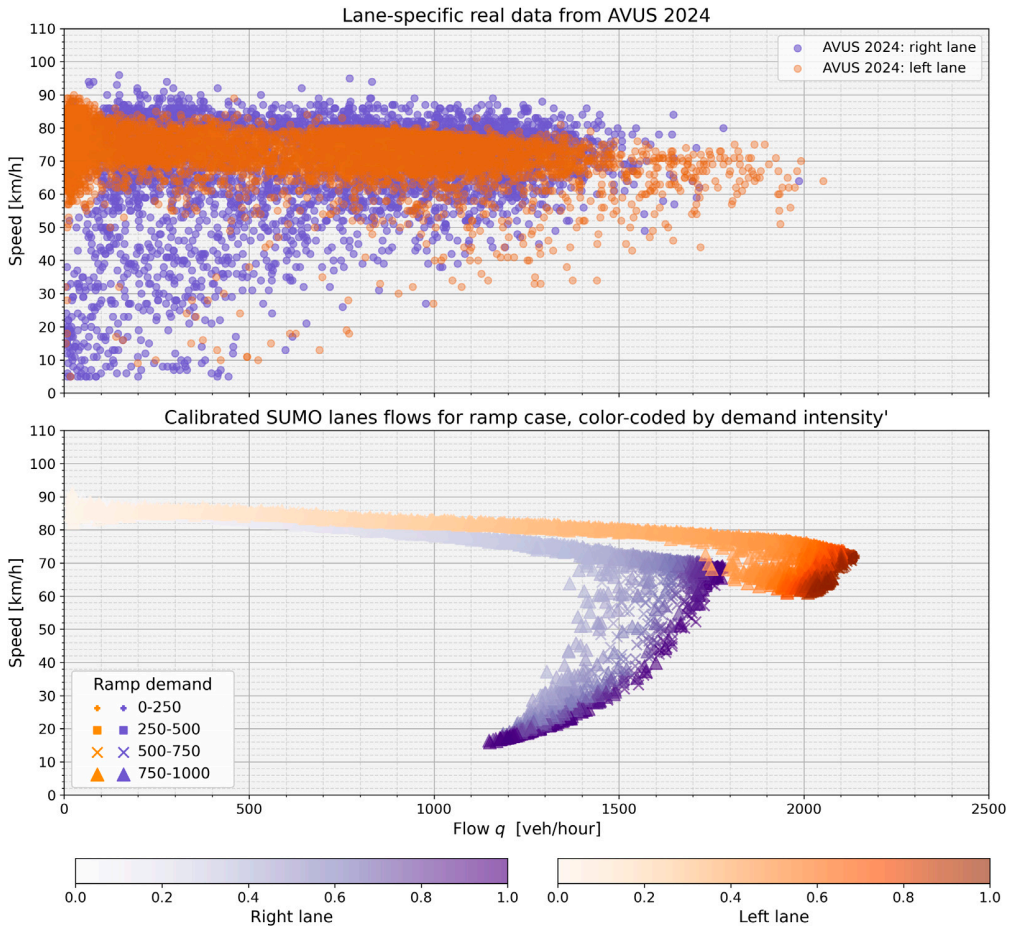

FIGURE 7
Lane-specific speed–flow relations for real data from 2024 (top panel) and SUMO's calibrated ramp scenario (bottom panel).

described in Section 2, this time varying the AV–MV share in 10%-intervals between the two vehicle types. The top panel in Figure 11a shows the time headway profiles for a 100% AV share, corresponding to the acceleration profile discussed in Figure 3b. The increasing headways for later-following AVs are clearly identifiable. In the bottom panel (Figure 11b), which depicts a 50–50 share, the headway increase is far less pronounced compared to the top panel.

Therefore, we introduce two additional estimators. In Equation 1, instead of using a fixed $\tau_{AV} = 1.6$ s, we derive τ_i for each share p_i from the numerical experiments as follows:

- `Max`: The black markers in Figure 11 denote the maximum time headway of each vehicle in the simulation run. The average of these maximum values serves as τ_i for each p_i in the estimator `max`.
- `Mean`: The point at which headways have stabilized after all ToCs are completed is marked by the vertical blue dashed line in Figure 11. Stabilization in this experiment is defined as the latest point after all headway peaks at which all vehicles' headways remain constant to within ±0.005 s for at least 15 s. The average of the time headways at this point serves as τ_i for each p_i in the estimator `mean`.

With these estimator-based τ-values, the original capacity Equation 1 is adjusted by replacing the fixed headway term in the denominator with $v\sum_i p_i \cdot \hat{\tau}_i(p_i)$, where $\hat{\tau}_i(p_i)$ is the empirically derived time headway for AVs as a function of their share p_i, while other vehicle types retain fixed values. The estimators `max` and `mean` provide these AV-specific headways based on the numerical experiments. Figure 10b presents the results of the calculations that account for ToCs by utilizing these estimators. Both trends exhibit a notable decline in estimated lane capacity compared to the ToC-ignorant estimation depicted in Figure 10a.

5 Results and discussion

Figure 12 presents the overall results obtained from the simulation study outlined in Section 4.1. First, we find that all maximum flows in the `AV00` share, ranging between

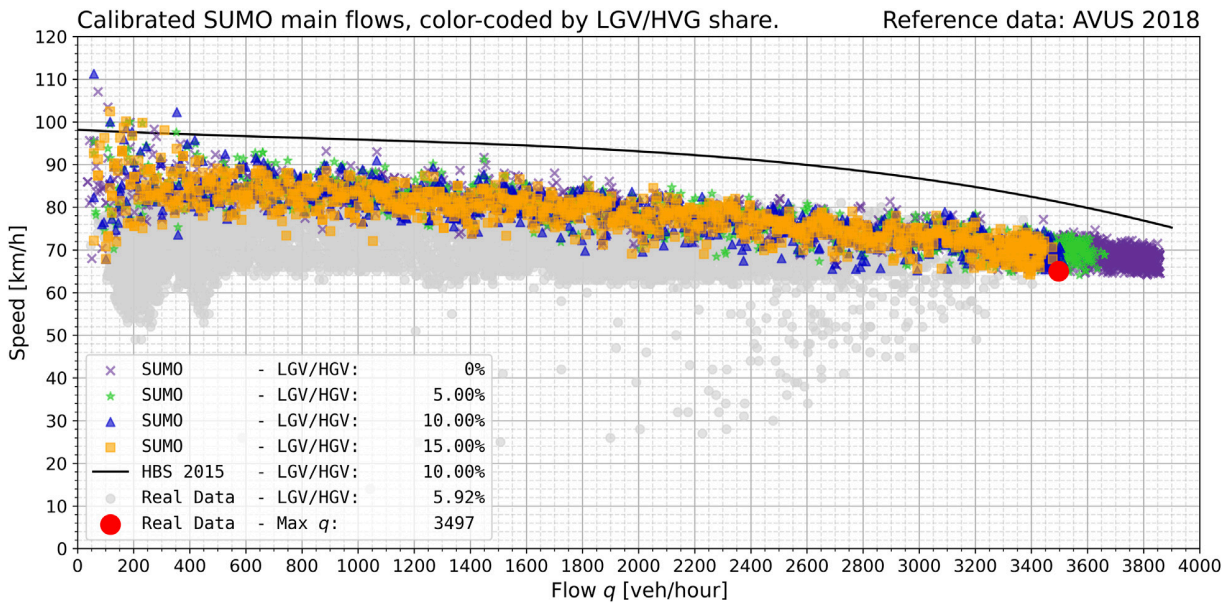

FIGURE 8
Speed–flow relations for the AVUS scenario based on the parametrization scheme from Table 2 comparing calibrated main flows considering LGV/HGV shares. Real-world detector data from AVUS 2018; color-coding by LGV/HGV share.

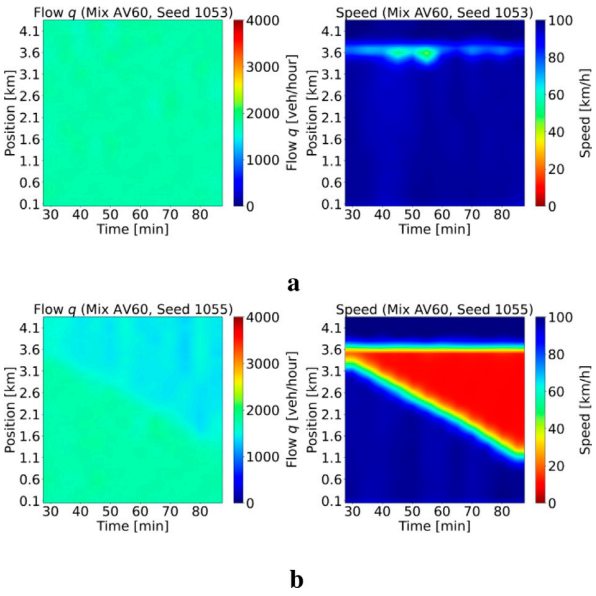

FIGURE 9
Spatiotemporal heatmaps of the ODD zone for speed and flow. **(a)** Valid run at max capacity: Mix AV60, seed 1053. **(b)** Invalid run: Mix AV60, seed 1055. The traffic breakdown is clearly visible in speed and flow.

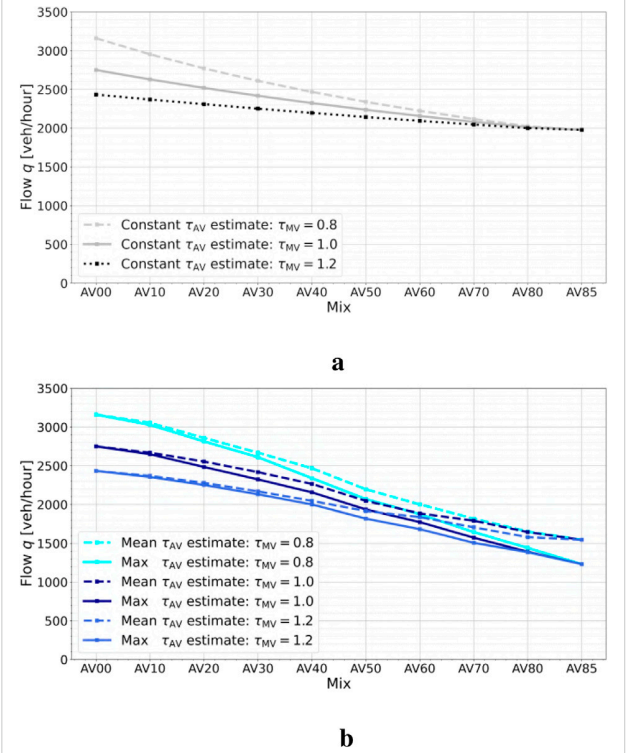

FIGURE 10
Estimated lane capacities based on Equation 1 for varied mean τ_MV-values. **(a)** shows capacity estimates assuming a constant τ_AV (no ToC consideration). **(b)** accounts for increasing τ_AV due to ToC effects (with ToC consideration).

3405 – 3611 veh/h, fall within the expected capacity range from the calibration, i.e., 3281 – 3647 veh/h. We further analyze these results in detail for the four scenarios, following the order in which they were previously defined:

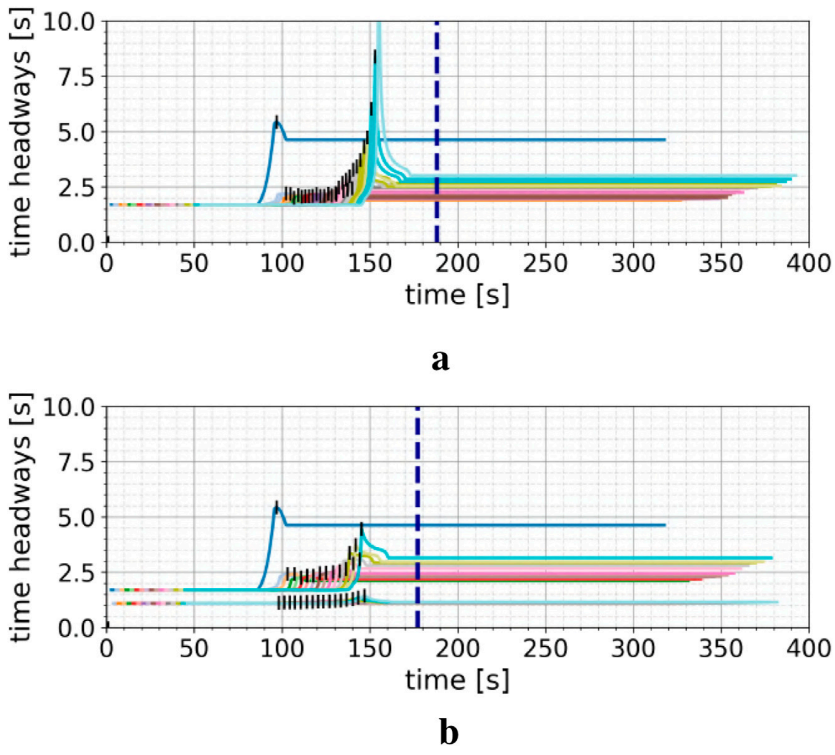

FIGURE 11
Time headways in a platoon experiment with 32 vehicles and varying AV–MV shares. Black markers indicate the maximum headway for each vehicle. The vertical blue dashed line marks the onset of system-wide headway stabilization (see text for criterion), after all ToCs are completed. **(a)** 100% AV vs. 0% MV share. **(b)** 50% AV vs. 50% MV share.

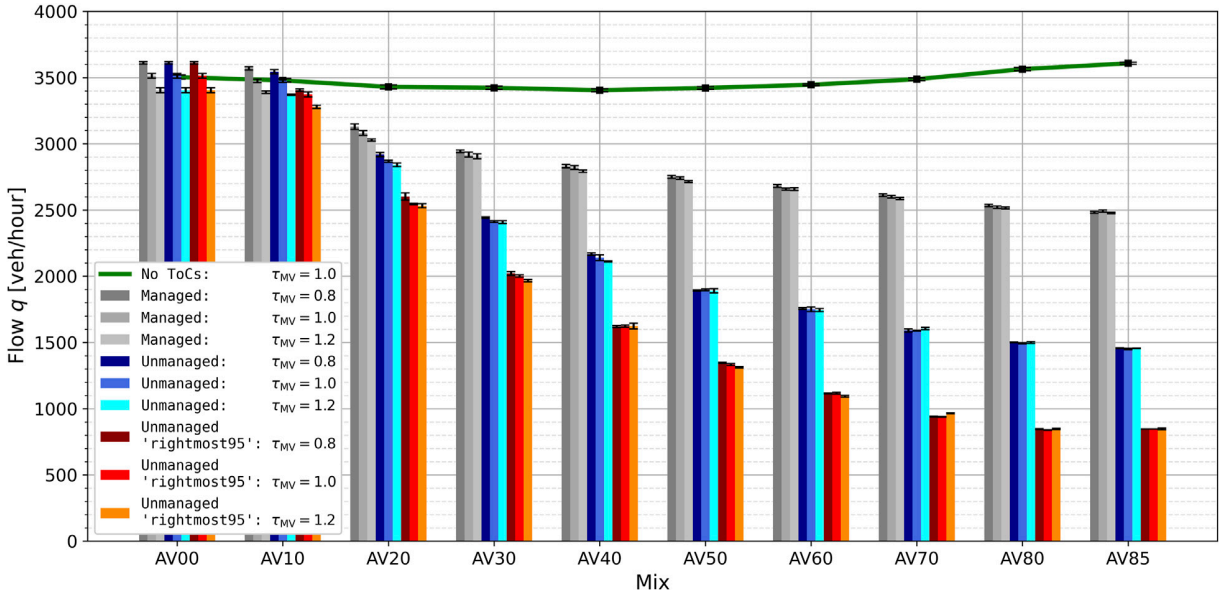

FIGURE 12
Maximum flow comparison across AV shares for the four scenarios: No ToCs (green line), Unmanaged ToCs (blue bars), Managed ToCs (gray bars), and Unmanaged "rightmost95" (red bars).

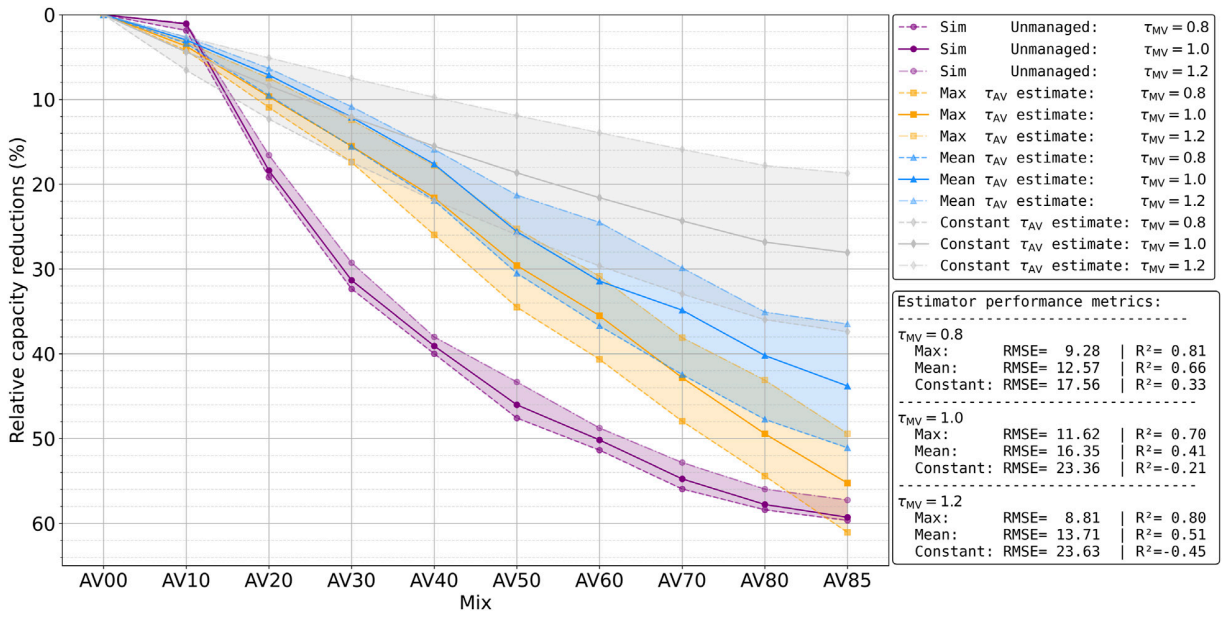

FIGURE 13
ToC-related relative percentage reductions of the capacity with varied mean τ_{MV}-values for the unmanaged scenario. For reference, the estimator "constant" (gray-colored), which does not account for ToCs, is included. The RMSE and R^2 metrics indicate estimation accuracy.

1. No ToCs: The results (green line) show that up to share AV40, maximum flows remain relatively stable, with a reduction of approximately 100 veh/h compared to AV00. Beyond AV50, flow values begin to increase again. This trend can be linked to a homogenization effect in traffic flow as AV shares grow, which is influenced by the AV parametrization—-specifically, the absence of sigma and a very small speedDev value of 0.01.
2. Unmanaged: In the case of entirely unmanaged ToCs (blue-colored bars), maximum flow decreases progressively from approximately 3500 veh/h at AV00 to around 1450 veh/h at AV85. Regarding the different τ_{MV} values, the results indicate high variations in maximum flow for AV00 and AV10. These variations become less pronounced as the AV share increases, starting around AV30.
3. Managed: When ToCs are managed within the ODD zone (gray-colored bars), maximum flows exhibit a similar decreasing trend but remain notably higher than in the unmanaged scenario. Flows decline from AV00 levels to approximately 2950 veh/h at AV85. As in the unmanaged case, variations related to τ_{MV} diminish with increasing AV shares, becoming noticeably less pronounced from AV30 onward.
4. Unmanaged "rightmost95": This scenario exhibits the lowest flow values across all AV shares (red-colored bars). A decline in maximum flow is already noticeable at AV20 and continues consistently as the AV share increases, reaching a minimum of 847 veh/h at AV85. In this scenario, capacity is inherently constrained because AVs are restricted to operating exclusively in the rightmost lane, leading to a disparate lane utilization with increasing AV share. Except for some LV and LGV vehicles traveling in the left lane, all HGVs and AVs remain on the right, thereby limiting capacity under unmanaged ToC conditions.

Overall, the results in Figure 12 show that in the unmanaged scenario, maximum flow declines significantly with increasing AV share. At AV85, the max flow is approximately 500 veh/h lower than in the managed case—indicating that ToC management measures could help alleviate, but not fully prevent, capacity losses.

Furthermore, to compare these simulation results with the theoretical lane capacities estimated in Section 4.2 and Figure 10c, we derive the relative percentage reductions in capacity across the increasing AV share. Figure 13 summarizes these reductions for the unmanaged scenario, differentiating between the estimators max and mean, while constant is included as a reference that ignores ToC effects. For each estimator, we report the root mean squared error (RMSE) and the coefficient of determination (R^2) to quantify the goodness of fit to the simulated capacity reductions—where lower RMSE and R^2 values closer to one indicate better agreement with the simulation data. The simulated results reveal a capacity loss of up to nearly 60% at AV85. We also make the following observations:

- While the τ_{MV} dependency is relatively small in the simulation results, it becomes increasingly important in the estimator outcomes.
- The notably poor performance of the constant estimator, including negative R^2 values in some cases, is expected since it does not capture capacity changes induced by ToCs.
- Both estimators, max and mean, although accounting for ToCs, notably underestimate the capacity reductions in the mid-range AV share (AV20–80). This can be attributed to the simplicity of the numerical experiments we conducted to derive the estimator

values. In particular, intensified vehicle interactions due to driver imperfections (parameter `sigma` σ) and speed factor variances are disregarded in these experiments. Additionally, the numerical experiments employ single-lane vehicle strings and uniform acceleration and deceleration parameters, omitting lane-changing interactions and parameter variability that are present in the two-lane simulation scenario (*cf.* Figure 5).

- Compared to the ToC-ignorant estimator `constant`, the other estimators perform notably better in predicting ToC-related capacity reductions, particularly `max`, which achieves the best RMSE and R^2 scores. At AV85 share, `max` matches best with the simulation results, as vehicle interaction effects with non-AVs have almost completely vanished (still 15% HGVs present), leading to minimal driving behavior variability that coincides with the numerical experiment setup, where all vehicles share the same parameterization.

In summation, the capacity reductions observed in the simulation might align only unsatisfactorily with theoretical estimates, as deviations occur in the mid-range AV shares due to the simplified assumptions of the estimators. This partial mismatch is also reflected in the RMSE and R^2 values, for which no established benchmarks exist in this context. Therefore, our assessment of estimator performance focuses on relative differences and qualitative trends within the observed results. However, the estimator `max` performs best in comparison to the simulation results, substantiating our suspicion that the maxima in time headway increments dominate ToC-related capacity effects. Nevertheless, the overall findings from Figures 12, 13 highlight the potential ToC effects on capacity reductions across various scenarios and parameter dependencies, in line with the stated expectations.

6 Conclusion

To investigate ToC-related capacity reductions, we conducted comprehensive simulation experiments with a calibrated two-lane highway scenario, as well as numerical experiments to estimate the large-scale impact of time headway increments during consecutive control transitions. Our main findings can be summarized as follows: (i) capacity reductions of up to 2000 veh/h, corresponding to approximately 60% loss, were observed in shares with near-full Level 3 automation but no traffic management coordination; (ii) ToC effects became notably impactful starting from a Level 3 share of 10% to 20%; (iii) a coordination of ToCs could mitigate losses by roughly 1000 veh/h or 30%; (iv) binding Level 3 operation to the rightmost lane resulted in the most severe reduction, with up to 2660 veh/h or 75% loss; and (v) maxima in time headway increments during ToCs emerge as the dominant factor contributing to these capacity effects.

Several relevant limitations should be acknowledged, as they may affect the applicability and interpretation of our findings. Recent research on data from Level 4 AVs reveals reduced time headways when MVs follow AVs (Jiao et al., 2024). Such effects, which might also apply to Level 3 systems, are not considered in this study. An additional aspect that has not yet been discussed is the impact of human response times for non-emergency ToCs. Throughout this investigation, the response time distribution was kept the same, at $\mu = 7$ s in all simulations, in line with our previous studies. More recent data from real-world tests presented by Pipkorn et al. (2023) indicate response times closer to 5 s, from which the authors infer that a lead time of 10 s, as specified in the R157, should be feasible for human drivers to take over in time. In our simulations, a few random sample reruns with these lower response times indicated approximately 10–20% higher capacities compared to the results presented here. A further limiting factor might be SUMO's ACC model, which is parametrized for full string stability and deployed here as a proxy for Level 3 automated vehicles, although experimental studies and theoretical work have demonstrated string instabilities in ACC-equipped platoons, as we also discussed in Alms and Wagner (2024). All these limitations could potentially affect traffic capacity, though their precise contribution cannot be reliably quantified at this stage.

Future work should therefore include the development of a more accurate Level 3 model in SUMO, for example, an ACC-based ALKS system, as well as a systematic investigation of how model assumptions and human response variability together affect traffic capacity. Another important direction is to further examine the effects of MRMs, which are relevant for failed Level 3 transitions and Level 4 automation, on overall traffic, especially if they are not managed properly.

Lastly, we would like to reflect on the broader capacity implications of AVs. Our overall vehicle parametrization inherently results in slightly reduced theoretical capacities—even without ToCs—due to the implementation of lower time headways for MVs and higher ones for AVs, which contrasts with assumptions commonly made in other studies. While experimental research has demonstrated counterbalancing effects at high AV shares, which our own simulations also imply, this effect is diminished in the context of Level 3 systems. Unlike Level 4 or CACC-equipped vehicles, Level 3 automation, in its current form, does not typically support the low time headways often assumed to contribute to capacity gains. However, practical capacity impacts at relevant market penetration rates between 10% and 20% are likely still many years away, leaving room for further technical and regulatory development of Level 3 systems. Yet, in combination with the ToC-related capacity constraints demonstrated in this study, we take a more cautious view and do not share the seemingly widespread optimism regarding beneficial capacity effects of AVs in the mid-term.

Data availability statement

The raw data supporting the conclusion of this article will be made available by the authors, without undue reservation.

Author contributions

RA: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Software, Validation, Visualization, Writing – original draft, Writing – review and editing. PW:

Conceptualization, Methodology, Supervision, Writing – review and editing.

Funding

The author(s) declare that no financial support was received for the research and/or publication of this article.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

Alms, R., Noulis, A., Mintsis, E., Lücken, L., and Wagner, P. (2022). Reinforcement learning-based traffic control: mitigating the adverse impacts of control transitions. *IEEE Open J. Intelligent Transp. Syst.* 3, 187–198. doi:10.1109/OJITS.2022.3158688

Alms, R., and Wagner, P. (2024). Control transitions in level 3 automation: safety implications in mixed-autonomy traffic. *Safety* 10, 1. doi:10.3390/safety10010001

Alvarez Lopez, P., Banse, A., Barthauer, M., Behrisch, M., Couéraud, B., Erdmann, J., et al. (2025). Simulation of urban mobility (SUMO). doi:10.5281/zenodo.14796685

Alvarez Lopez, P., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.-P., Hilbrich, R., et al. (2018). "Microscopic traffic simulation using SUMO," in The 21st IEEE International Conference on Intelligent Transportation systems *(IEEE)*, 2575–2582.

BASt (2021). *Road Traffic Census Report 2021 - Ergebnisbericht der Straßenverkehrszählung 2021*. Tech. rep. *Fed. Highw. Res. Inst. (BASt)*.

BASt (2025). *Straßenverkehrszählung – verkehrsdatenbank*.

Bolovinou, A., Anagnostopoulou, C., Roungas, V., Amditis, A., González, R. B., Coello, L. T., et al. (2023). HI-DRIVE Deliverable D3.1/Use cases definition and description. *Tech. Rep. Eur. Comm.*

Brilon, W., and Geistefeldt, J. (2010). Überprüfung der Bemessungswerte des HBS für Autobahnabschnitte außerhalb der Knotenpunkte*, vol. 1033 of* Forschung Straßenbau und Straßenverkehrstechnik. Bremerhaven, Germany: Wirtschaftsverl. NW Verl. für neue Wissenschaft.

British Standards Institution (2020). *PAS 1883:2020 - operational design domain (ODD) taxonomy for automated driving systems (ADS) – specification*. London, UK: British Standards Institution.

Calvert, S., and van Arem, B. (2020). A generic multi-level framework for microscopic traffic simulation with automated vehicles in mixed traffic. *Transp. Res. Part C Emerg. Technol.* 110, 291–311. doi:10.1016/j.trc.2019.11.019

Digitale Plattform Stadtverkehr Berlin (2024). *Verkehrsdetektion berlin*.

FGSV (2015). *Handbuch für die Bemessung von Straßenverkehrsanlagen: HBS 2015. No. FGSV 299 B in FGSV W1 - Wissensdokumente* (Cologne, Germany: FGSV-Verl). 2015 Edn.

Friedrich, B. (2016). *The effect of autonomous vehicles on traffic*. Berlin, Heidelberg: Springer Berlin Heidelberg, 317–334. doi:10.1007/978-3-662-48847-8_16

Geistefeldt, J., Giuliani, S., Busch, F., Schendzielorz, T., Haug, A., Vortisch, P., et al. (2017). "HBS-conform simulation of freeway traffic flow," in *vol. 279 of Berichte der Bundesanstalt für Straßen- und Verkehrswesen, Reihe V: Verkehrstechnik* (Germany: Federal Highway Research Institute).

Jiao, Y., Li, G., Calvert, S. C., van Cranenburgh, S., and van Lint, H. (2024). Beyond behavioural change: investigating alternative explanations for shorter time headways when human drivers follow automated vehicles. *Transp. Res. Part C Emerg. Technol.* 164, 104673. doi:10.1016/j.trc.2024.104673

Kohanpour, E., Davoodi, S. R., and Shaaban, K. (2025). Trends in autonomous vehicle performance: a comprehensive study of disengagements and mileage. *Future Transp.* 5, 38. doi:10.3390/futuretransp5020038

Liu, Q., Gao, C., Wang, H., Cai, Y., Chen, L., and Lv, C. (2024). Learning from trajectories: how heterogeneous CACC platoons affect the traffic flow in highway merging area. *IEEE Trans. Veh. Technol.* 73, 16212–16224. doi:10.1109/TVT.2024.3419143

Lücken, L., Mintsis, E., Porfyri, K., Alms, R., Flötteröd, Y.-P., and Koutras, D. (2019). "From automated to manual - modeling control transitions with SUMO,"in SUMO user conference 2019. Editors M. Weber, L. Bieker-Walz, R. Hilbrich, and M. Behrisch (Stockport, United Kingdom: EPiC Series in Computing), 62, 124–144. doi:10.29007/sfgk

Maerivoet, S., Akkermans, L., Carlier, K., Flötteröd, Y.-P., Lücken, L., Alms, R., et al. (2019). TransAID Deliverable 4.2 - preliminary simulation and assessment of enhanced traffic management measures. *Tech. Rep. Eur. Comm.*

McDonald, A. D., Alambeigi, H., Engström, J., Markkula, G., Vogelpohl, T., Dunne, J., et al. (2019). Toward computational simulations of behavior during automated driving takeovers: a review of the empirical and modeling literature. *Hum. Factors* 61, 642–688. doi:10.1177/0018720819829572

Mercedes-Benz Group (2024). Mercedes-Benz is approved for 95 km/h Level 3 autonomous driving in Germany. *Tech. Rep. Mercedes Benz Group*.

Mintsis, E., Koutras, D., Porfyri, K., Mitsakis, E., Lücken, L., Erdmann, J., et al. (2019). TransAID Deliverable 3.1 - modelling, simulation and assessment of vehicle automations and automated vehicles' driver behaviour in mixed traffic. *Tech. Rep. Eur. Comm.*

Park, J. E., Byun, W., Kim, Y., Ahn, H., and Shin, D. K. (2021). The impact of automated vehicles on traffic flow and road capacity on urban road networks. *J. Adv. Transp.* 2021, 1–10. doi:10.1155/2021/8404951

Pipkorn, L., Tivesten, E., Flannagan, C., and Dozza, M. (2023). Driver response to take-over requests in real traffic. *IEEE Trans. Human Mach. Syst.* 53, 823–833. doi:10.1109/THMS.2023.3304003

Rummel, J. (2017). *Replication of the hbs autobahn with sumo*. Berichte aus dem DLR-Institut für Verkehrssystemtechnik, Berlin, Germany: Institute of Transportation Systems, German Aerospace Center 171–178.

SAE International (2021). SAE international recommended practice: taxonomy and definitions for terms related to driving automation systems for on-road motor vehicles. *SAE Int.* doi:10.4271/J3016_202104

Sauvaget, J.-L., Dakil, M., Griffon, T., Anagnostopoulou, C., Bolovinou, A., Sintonen, H., et al. (2023). HI-DRIVE deliverable D5.1/descriptions of "operations". *Tech. Rep. Eur. Comm.*

Schrader, M. (2024). *Calibrating traffic microsimulation for optimization of intelligent transportation systems using roadside radar*. Tuscaloosa, AL: The University of Alabama. Ph.d. thesis.

Schulte-Tigges, J., Matheis, D., Reke, M., Walter, T., and Kaszner, D. (2023). "Demonstrating a V2X enabled system for transition of control and minimum risk manoeuvre when leaving the operational design domain," in *HCI in mobility, transport, and automotive systems*. Editor H. Krömker (Cham: Springer Nature Switzerland), 200–210. doi:10.1007/978-3-031-35678-0_12

Shuttleworth, J. (2019). SAE updates J3016 automated-driving graphic. *Mobil. Eng.*

TransAID (2021). Transition areas for infrastructure-assisted driving (TransAID). *H2020 Res. Proj. 723390, Eur. Comm.*

UNECE (2023). Addendum 156 – UN regulation No. 157 - amendment 4 - uniform provisions concerning the approval of vehicles with regard to automated lane keeping systems. *Tech. Rep.* United Nations Economic commission for Europe.

Van Aerde, M. (1995). "A single regime speed-flow-density relationship for freeways and arterials," in *74th annual meeting of the transportation research board*. Preprint paper no. 950802.

Van Lint, J., and Calvert, S. C. (2018). A generic multi-level framework for microscopic traffic simulation—theory and an example case in modelling driver distraction. *Transp. Res. Part B Methodol.* 117, 63–86. doi:10.1016/j.trb.2018.08.009

Wagner, P. (2012). Analyzing fluctuations in car-following. *Transp. Res. Part B Methodol.* 46, 1384–1392. doi:10.1016/j.trb.2012.06.007

Wang, C., Ren, W., Xu, C., Zheng, N., Peng, C., and Tong, H. (2024). Exploring the impact of conditionally automated driving vehicles transferring control to human drivers on the stability of heterogeneous traffic flow. *IEEE Trans. Intel. Veh.*, 1–17. doi:10.1109/TIV.2024.3419789

Wang, C., Xu, C., Peng, C., Tong, H., Ren, W., and and, Y. J. (2025a). Predicting the duration of reduced driver performance during the automated driving takeover process. *J. Intel. Transp. Syst.* 29, 218–233. doi:10.1080/15472450.2024.2307029

Wang, C., Xu, C., Shao, Y., Zheng, N., Peng, C., Tong, H., et al. (2025b). Identifying factors affecting driver takeover time and crash risk during the automated driving takeover process. *J. Transp. Saf. Secur.* 0, 1–26. doi:10.1080/19439962.2025.2450695

Ward, T. (2024). "Areas of improvement for autonomous vehicles: a machine learning analysis of disengagement reports," in 2024 International Conference on Electrics and Computer (INTCEC). ArXiv:2408.00051.

Xiao, L., Wang, M., and van Arem, B. (2017). Realistic car-following models for microscopic simulation of adaptive and cooperative adaptive cruise control vehicles. *Transp. Res. Rec.* 2623, 1–9. doi:10.3141/2623-01

Chapter synthesis

Building on earlier project work, notably TransAID, which provided the initial impetus and boundary conditions for this research, the problem frame was refined and the overarching research question emerged. The three central aspects, efficiency, safety and capacity, were then examined in depth in Papers I–III as separate studies: RL-based ToC management for efficiency, ODD-constrained safety effects, and ToC-induced capacity impacts, addressing RQ 1, RQ 2 and RQ 3 respectively. Together these publications establish the empirical context and modelling choices for the complementary analyses and integration developed in the subsequent discussion Chapter 4.

4 Discussion

This chapter contextualises the findings of Papers I–III with complementary results to provide a fuller basis for addressing the three research questions. The added analyses address specific shortcomings and evidential gaps. For Paper I, further RL experiments were undertaken to probe policy sensitivity and the efficiency versus automated-distance trade-off. For Paper II, safety results were complemented by applying the motorway scenario inspired by use case 5.1 from Paper III to examine a different ODD boundary case (the No-AD zone). For Paper III, supplementary measurements from the main two-lane simulations and from the single-lane numerical experiments were included to examine time headway estimator bias and the role of lane changes. The discussion treats these shortcomings explicitly in light of the additional results and clarifies which conclusions are robust and which depend on modelling choices or parameterisation.

4.1 Methodological Reflections on RL-based ToC Management

The study in Paper I showed that reinforcement learning can be used to mitigate the adverse effects of system-initiated control transitions on traffic flow. By contrasting unmanaged scenarios, heuristic approaches, and RL-based control, it illustrated both the promise and the shortcomings of such methods. Although the original results suggested that RL controllers may outperform baseline strategies in balancing traffic efficiency and automated driving distance, further reflection is needed regarding the robustness of these outcomes. In the following, additional experiments with updated RL libraries and alternative algorithms are discussed to assess the sensitivity of the results to methodological choices and reward design.

4.1.1 From SB2 to SB3: Progress and Local Minima

To examine the robustness of the results from Paper I, the experiments were retrained with an updated version of the Stable Baselines library [66] (SB2 from 2021 and

SB3 from 2025), using the original TD3 algorithm on both versions and, in addition, a PPO run on SB3 (see Section 4.1.3). Figure 4.1 presents these results: the top panel shows the development of episode rewards, and the bottom panel compares the resulting traffic metrics, namely average travel time and CAV distance, against the baseline cases from Paper I.

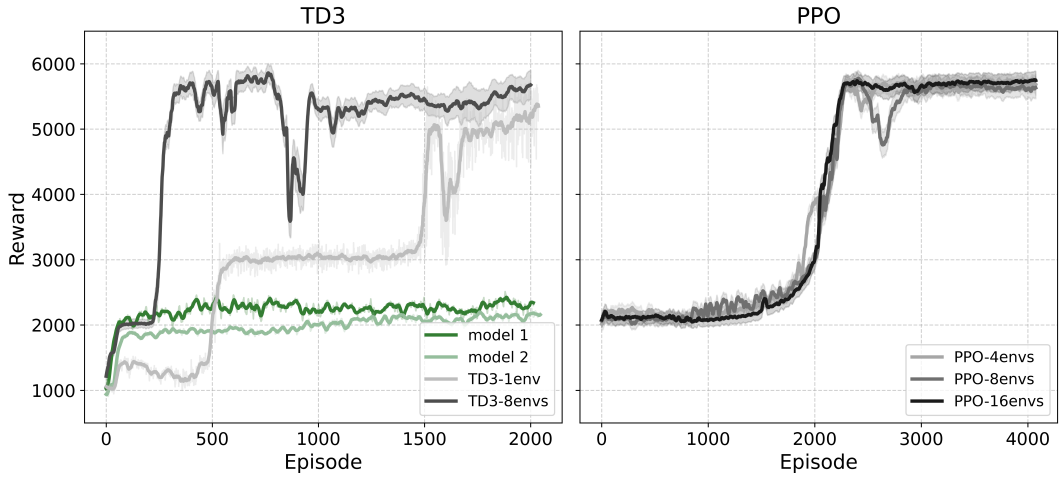

(a) Training reward development for TD3 and PPO.

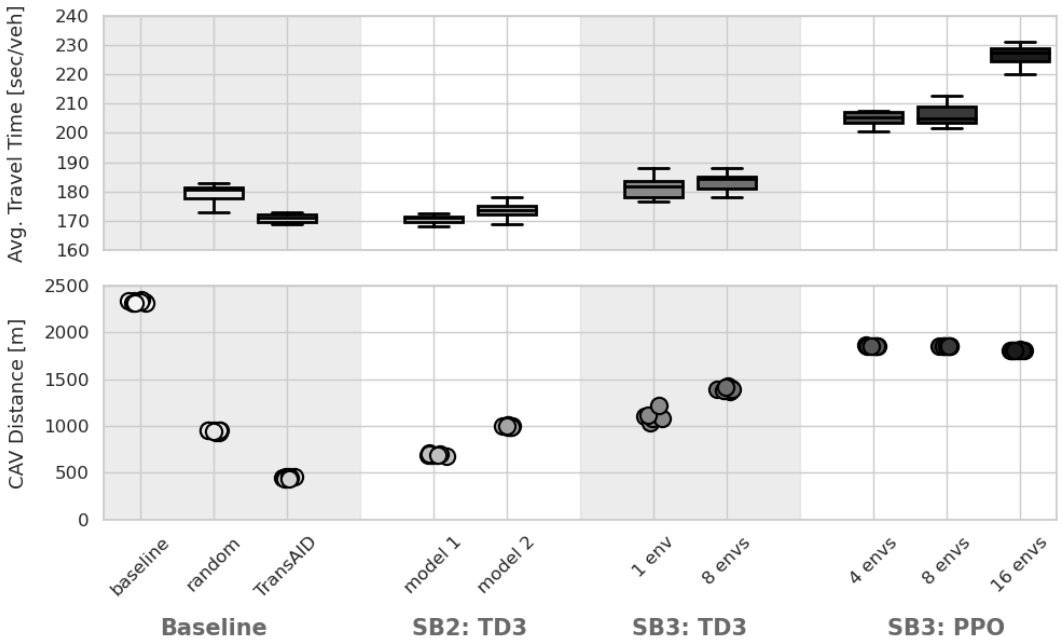

(b) Average travel time and CAV distance across methods.

Figure 4.1: Training and traffic performance for TD3 and PPO (TD3: SB2 and SB3, PPO: SB3). **Panel (a):** Episode rewards under different training environment configurations. For TD3, SB3 attains higher rewards than with SB2. **Panel (b):** Evaluation metrics (average travel time and CAV distance). Higher rewards do not yield better efficiency: under SB3, policies shift towards maximising automated driving continuation at the cost of higher average travel time, whereas TD3 with SB2 (models 1 and 2) remains more balanced.

The training curves for model 1 and model 2 (cf. Figure 4.1a, TD3) show that the SB2 runs increased in reward initially but then levelled off early, indicating limited further learning. In contrast, the SB3 runs for TD3 continued to improve throughout training, even from initially lower reward levels, and reached noticeably higher final reward values, reflecting greater effectiveness of the updated implementation. The bottom panel of Figure 4.1 reveals the consequences of this difference: SB2, despite its restricted learning progress, produced a comparatively balanced outcome between travel efficiency and automated driving distance, whereas SB3, by exploiting the reward function more fully, shifted strongly towards maximising automated driving distance at the cost of higher average travel times, thereby exposing a flaw in the reward design, which is discussed more detailed in the following.

4.1.2 Reward Imbalance and Traffic Efficiency Trade-offs

Based on the idea of discretising the road into 14 cells (7 per lane) to which the RL agent could send a ToR to a group of CAVs, the training algorithm aimed at maximising the reward by finding a compromise between keeping a decent overall speed level (reward: R_{speed}) in the scenario while also allowing AVs to continue their automated driving mode (reward: R_{ToR}) as long as feasible before entering the No-AD zone. This idea, represented in equation (1) in Paper I, can be re-written as:

$$R \;=\; R_{\text{speed}} \;+\; R_{\text{ToR}}. \tag{4.1}$$

The flaw in this reward, as indicated by the results from Section 4.1.1, is that the term R_{ToR} is parametrised by the weights w_{ToR} and p_{ToR}[11] to favour delaying ToRs excessively. With $p_{\text{ToR}} = 10$ and $w_{\text{ToR}} = 100$, issuing a ToR in downstream (non-final) cells yields a large *action-dependent* credit (scaled by w_{ToR}), while the penalty applied in the final cells (scaled by p_{ToR}) is comparatively weak. Consequently, the agent learns to postpone ToRs to late (penultimate) cells, thereby increasing automated-mode distance at the expense of average travel time.

[11] w_{ToR} and p_{ToR} (denoted π_{ToR} in Paper I) are free parameters for reward shaping.

More precisely, with $i \in \{0, \ldots, 13\}$ indexing the cells, n_i denotes the number of automated vehicles (AVs) in cell i, and $a_i \in \{0, 1\}$ indicates whether a ToR is issued to cell i (1 if issued, 0 otherwise). The parameters p_{ToR} and w_{ToR} enter the formulation as follows:

$$R_{\text{ToR}} = \sum_{i=0}^{13} a_i \, n_i \, r_i \; - \; p_{\text{ToR}} \sum_{i \in \{12,13\}} n_i, \tag{4.2}$$

$$r_i = \begin{cases} w_{\text{ToR}} \, \dfrac{\lfloor i/2 \rfloor}{6}, & i \notin \{12, 13\}, \\ p_{\text{ToR}}, & i \in \{12, 13\}, \end{cases} \qquad w_{\text{ToR}} = 10, \;\; p_{\text{ToR}} = 100.$$

The practical effect of this parameterisation was that SB3 agents, by exploiting the reward term more effectively, maximised automated driving distance at the cost of increased average travel times. In fact, in some cases the SB3 policies were outperformed in terms of traffic efficiency by the random and TransAID baselines. As Figure 4.1 illustrates, SB2 plateaued in a local minimum yet delivered a comparatively balanced compromise, whereas SB3 reached higher reward values while exposing the imbalance in the reward function. The lesson from this comparison is that reward design as well as its parametrization (p_{ToR} and w_{ToR}) must adequately reflect the intended trade-off between traffic efficiency and automated driving comfort.

4.1.3 Additional PPO Experiments

To verify whether the poorer traffic outcomes with TD3 were due to algorithmic limitations or the reward imbalance identified above, additional experiments were conducted with PPO. In contrast to TD3's continuous action space, PPO employs a discrete design in which the agent directly choses whether to issue a ToR (1) or not (0). Training was extended to 4000 episodes and carried out with multiple parallel environments (4, 8, or 16).

As shown in Figure 4.1, PPO eventually reached reward levels comparable to TD3 with SB3, but only after a longer training horizon[12]. Learning progressed more smoothly and stably, reflecting the robustness of the on-policy approach. However, this stability did not translate into better traffic performance: average travel times continued to rise, and CAVs remained longer in automated mode at the expense of overall efficiency. These results support the inference that the underlying issue lies in the reward imbalance rather than in the choice of the RL algorithm.

[12]The apparent sample-efficiency gap between the algorithms is explained by their fundamentally different data usage. TD3 in the experiments collected transitions from a single environment and reused them many times via a replay buffer (off-policy), whereas PPO relied on parallel environments to generate a much larger number of transitions, each consumed only once (on-policy).

4.1.4 Implications for RL-based ToC Management

Across the experiments, the move from SB2 to SB3 increased training stability and, for the same TD3 controller, raised the achieved reward levels relative to those reported in Paper I. Yet, as Figure 4.1 shows, these gains did not translate into better traffic outcomes: the original SB2 run, despite plateauing early, yielded a comparatively balanced compromise, whereas the SB3 retraining achieved higher rewards while notably worsening average travel times (for a fixed 5 km trip in use case 5.1: 170 s $\approx$ 106 km/h; 227 s $\approx$ 79 km/h – about 25% lower network speed). Additional PPO runs verify that this pattern is not algorithm-specific. Taken together, these results expose a gap between the internal optimisation target in the MDP formalisation (reward) and external multiple objectives (traffic efficiency vs. CAV distance) in the RL experiment.

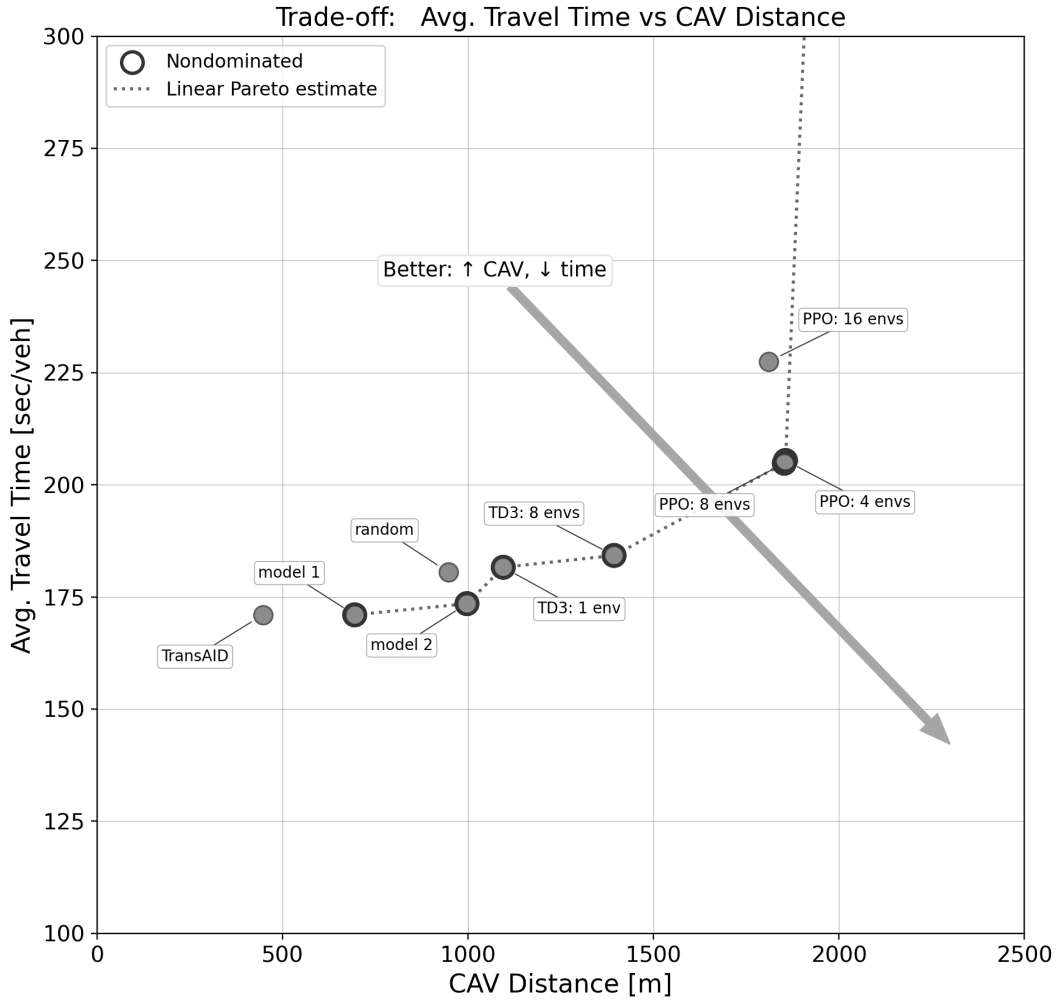

Figure 4.2: Trade-off between average travel time and CAV distance in use case 5.1 (data from Figure 4.1). Outlined points are non-dominated within the sample; the dotted line is a linear estimate of a potential Pareto frontier. The arrow indicates the desirable direction (lower travel time, longer automated distance).

Figure 4.2 illustrates this trade-off by plotting average travel time against CAV distance. The outlined markers indicate the non-dominated points, i.e. a candidate Pareto-efficient set in the multi-objective RL sense [67, 51]. The dotted line is a linear fit to this potential frontier within the explored [13]. Further progress with RL-based ToC management in this scenario depends on (i) reward formulations that explicitly capture the intended trade-off between traffic efficiency and automated driving continuation and (ii) multi-objective RL methods that operate on sets of non-dominated policies rather than a single scalar reward (cf. [67, 51]).

4.2 Safety Implications of ODD-constrained ToCs

Paper II investigates ToC-related traffic safety in greater depth, as Paper I raised potential concerns without further quantification. It examines a use case different from Paper I, motivated by Level 3 systems certified under the 2021 version of UNECE Regulation No. 157 [79], which, at the time, limited the ODD to 60 km/h. First, Section 4.2.1 addresses the rationale for using SSMs in simulation, notes SUMO-specific caveats, and clarifies the evaluation approach used in the subsequent analysis. Section 4.2.2 then presents complementary safety results from additional simulations conducted for Paper III, which also align with the use case 5.1 concept from Paper I. The 60 km/h ODD-limited use case from Paper II is revisited in Section 4.2.3.

4.2.1 SSMs in SUMO

Assessing traffic safety in simulations, particularly in SUMO, is challenging because crashes are rare by design and model assumptions limit absolute claims. Therefore, SSMs are used in this research to reveal and compare safety-relevant changes that do not culminate in collisions. Paper II employs a pair of indicators (TTC & MDRAC), following [94], which shows that combining time- and deceleration-based measures (e.g., TTC & DRAC) provides complementary, non-redundant signals and broader coverage of the investigated conflict types than single measures alone. SUMO's driver models do not explicitly represent human perception–reaction time (PRT), which can make DRAC overly optimistic, particularly when reactive controllers induce oscillatory braking. To address this, MDRAC [34], which embeds PRT, was adopted for greater sensitivity to imminent conflicts. Most SSMs, including TTC and DRAC, are typically evaluated with fixed thresholds. In SUMO these thresholds are not empirically validated, so exceedance counts have limited interpretive value. Even

[13] The baseline point lies outside the plotting window and is intentionally excluded, as it corresponds to a post–breakdown traffic regime that is not directly comparable to the other operating points. For context, the dotted line is extrapolated beyond the visible range toward that point.

alternative Level 2 takeover metrics, such as the margin between time to control and a safe time budget [58], depend on provisional thresholds. Consequently, the safety analysis here prioritises distributional shifts in the well-established TTC, supported by the MDRAC metric, to reveal systematic changes rather than relying on critical events defined by fixed thresholds.

4.2.2 Complementary Safety Results

Prior safety results for the ODD-limited use case 5.1 (cf. Figure 3.2) already indicated ToC-induced deterioration in a stable traffic stream. Because those analyses covered only lower AV shares (15 to 40%), simulations from Paper III were reproduced to cover the full range of synthetic AV shares. Figure 4.3 shows the resulting per-hour SSM event counts, disaggregated into inter- and intra-vehicle interactions. The unmanaged case yields markedly higher event rates than the managed case across traffic mixes, with the largest differences in inter-vehicle interactions at intermediate AV shares. Interestingly, the pattern across traffic mixes peaks around 20% AVs and attenuates from about 40% upwards, indicating the worst safety effects at intermediate mixed-autonomy compositions, with both TTC and MDRAC exhibiting the same qualitative shape across mixes. This is supported by inter-vehicle event rates exceeding intra-vehicle event rates.

4.2.3 Safety under a 60 km/h ODD Limit

In the main 60 km/h ODD-limited use case of Paper II, safety signals deteriorate as the AV share increases and v/c ratios approach 100%. The tabular summaries show higher hourly rates of critical SSM events (rising from near-zero to about 7 to 9 events per hour in the most affected mixes), while the heatmaps reveal the distributional shifts more clearly, with a consistent pattern across both SSMs. The detrimental effects vanish when the ToC preparation controller is parametrised to avoid active deceleration by assuming a headway increase from the lead vehicle's acceleration. This finding contrasts with the main case, indicating that ToC-induced deceleration at the ODD boundary disrupts traffic flow and, through interactions with accelerating vehicles, is the primary driver of the observed safety deterioration. Deploying a string-unstable ACC parameterisation, as reported in the original model introduction by [47], markedly amplifies the detrimental effects, which is evident from broader and stronger adverse shifts in the respective heatmaps. These stability implications, amongst other aspects, are discussed in the following.

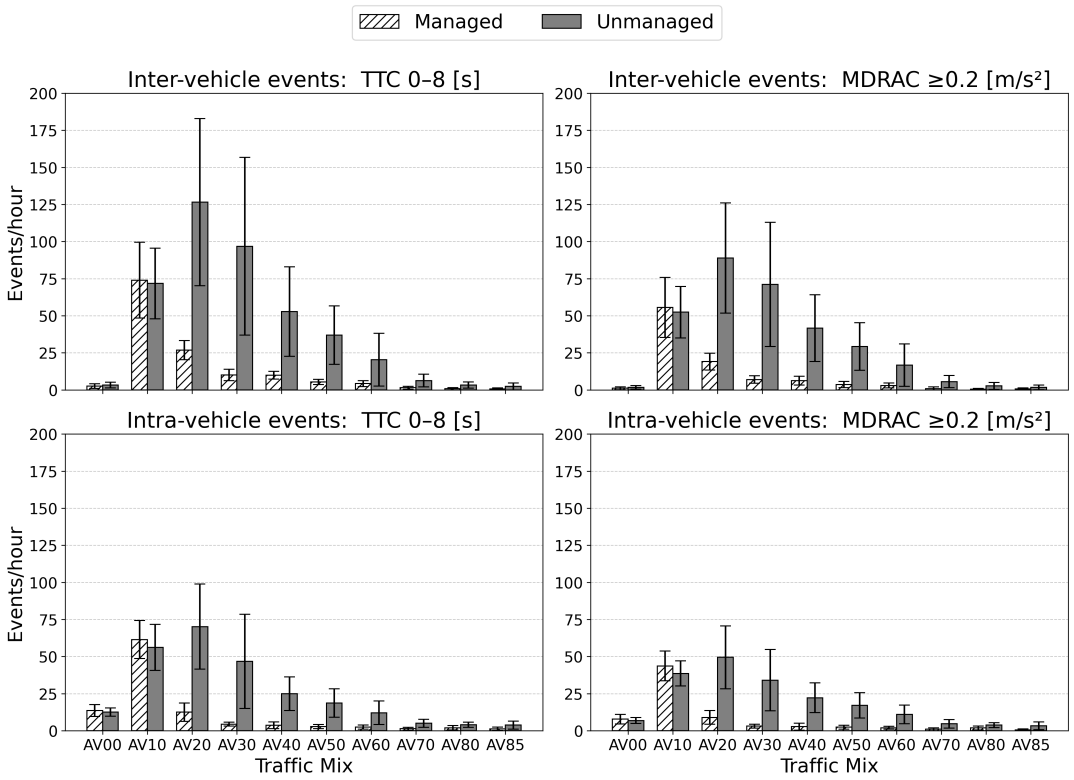

Figure 4.3: Average number of SSM events per simulation hour across traffic mixes. **Top panels:** inter-vehicle conflicts (between different vehicle types). **Bottom panels:** intra-vehicle conflicts (between the same vehicle type). **Left**: TTC 0–8 [s]. **Right**: MDRAC $\geq 0.2\,[\mathrm{m/s^2}]$. Results are shown for managed and unmanaged scenarios from Paper III. For comparability, both scenarios are evaluated at the demand level of the unmanaged case at 100% v/c ratio, since higher throughput in the managed case would otherwise bias the number of events.

4.2.4 Safety Consequences of ToCs at ODD Boundaries

Two different ODD-boundary settings are analysed through distributional TTC and MDRAC: (i) the 60 km/h ODD-limited case and (ii) the No-AD zone ODD-end case. Across both settings, safety signals worsen as AV share rises and as v/c ratio approaches 100%, while managed operation outperforms unmanaged operation. With a constant speed boundary (ii), differences are largest at intermediate AV shares, whereas with an increasing speed boundary (i), impacts are worst at maximum AV share. In (i), a ToC preparation without active deceleration removes this exacerbation. Inter-vehicle events dominate over intra-vehicle events, and effects peak in mixed traffic near 20% before attenuating from about 40% upwards.

Stability further modulates these outcomes: with a string-unstable ACC parameterisation, safety deterioration intensifies. Two notable studies considered here investigate stability effects of Level 3 conditional AVs specifically: In [39], the authors base conditional AV control on a model calibrated to SUMO simulations encompassing ToCs using the CACC model (cf. [47]) and then analytically show that

the resulting controller is string-stable. Given that the CACC model is inherently parametrised to be string-stable in SUMO, their stability-versus-penetration results therefore hold only under these string-stable controller assumptions. Conversely, [85] show that stability worsens with more ToCs. For any fixed takeover rate ($p_{\text{tor}} < 1$), it improves with higher Level 3 penetration because the automated, non-delayed dynamics compensate for the destabilising effects of the delayed post-takeover mode. Nevertheless, those findings do not translate easily to SUMO's default calibration, as its driver and ACC models are tuned near string stability and to collision avoidance. To study stability effects reliably, instability must be induced deliberately, for example via parameter choices outside stable ranges or by model extensions, as proposed by [41]. Thus, the stability impact on safety reported here is strongly parameterisation sensitive.

Taken together, the detrimental outcomes induced by disruptive, ODD-constrained ToCs arise predominantly from adverse vehicle interactions that propagate upstream and can lead to safety-critical events. However, quantification of these effects is bounded by modelling limits and parameterisation. Human PRT is not explicitly modelled in SUMO and MDRAC only proxies it. A calibrated crash-severity measure is also absent, so comparisons are necessarily relative to a baseline without ToCs rather than predictive.

4.3 Quantifying Capacity under ToC Constraints

Whilst Paper I only briefly discussed capacity implications of ToCs, Paper II considered reduced capacities in a limited set of baseline measurements. Paper III presented an in-depth quantification of ToC-induced capacity effects for an ODD-limited motorway scenario. Therefore, Sections 4.3.1 to 4.3.3 examine interdependencies of ToC-related factors that affect capacity.

4.3.1 Time Headway Distributions in the ODD-limited Motorway Scenario

Figure 4.4 presents the time headway distributions across AV shares for unmanaged and managed cases in the two-lane ODD-limited motorway simulations of Paper III. As AV penetration rises, both the mean and the right-hand tail increase, most markedly without management, whilst management moderates both the mean and spread at the same share. These patterns corroborate the capacity-relevant differences reported in Paper III. Table 4.1 lists the respective mean values for the unmanaged case.

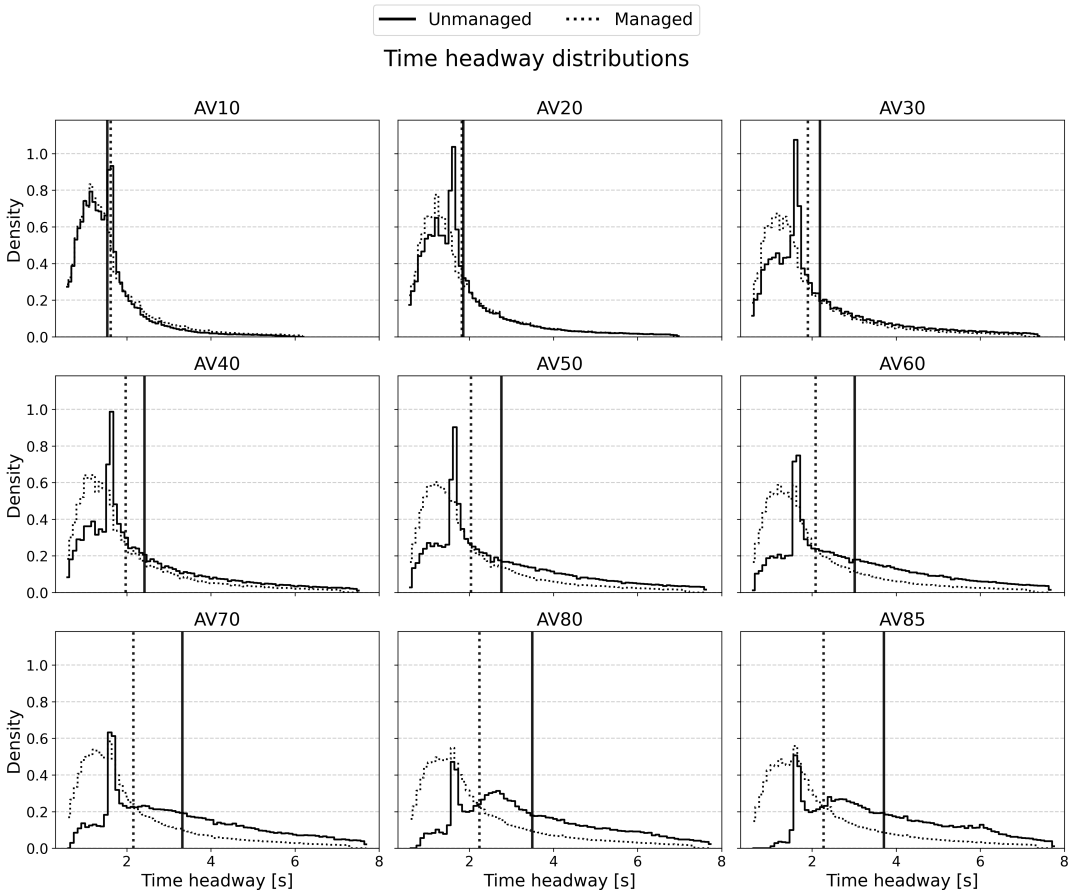

Figure 4.4: Time headway distributions by AV share (AV10–AV85) for unmanaged and managed cases from Paper III simulations. Vertical lines mark the mean of each distribution.

4.3.2 From Numerical Experiments to Capacity Estimates

Paper III linked ToC-induced maxima in time headways from a controlled single-lane 32-vehicle string to network capacity. Here, "numerical experiments" refers to controlled computations on the single-lane string used to elicit peak AV time headways at a fixed ODD boundary. To isolate ToC effects, the micro experiment was designed without lane changes, lateral interactions and behavioural heterogeneity, and the spatial ordering of AVs and MVs was prescribed. Therefore, to check for bias due to the equal AV—MV spacing used in Paper III, the experiment was repeated with two alternative patterns within the platoon, random placement and clustered AVs. Figure 4.5 reports the resulting peak headway estimator $\hat{\tau}_{\mathrm{AV}}^{\max}(p_i)$ across AV shares for all three patterns.

The results from Figure 4.5 reveal that all patterns yield fairly similar peak time headways, although the gap between patterns slightly widens as AV shares increases[14]. In Paper III, the equal-spacing estimates were propagated through the

[14]Within each placement pattern, variability across permutations is reflected in the violin widths. This indicates that the intra-platoon distribution indeed affects time headway maxima, but only to

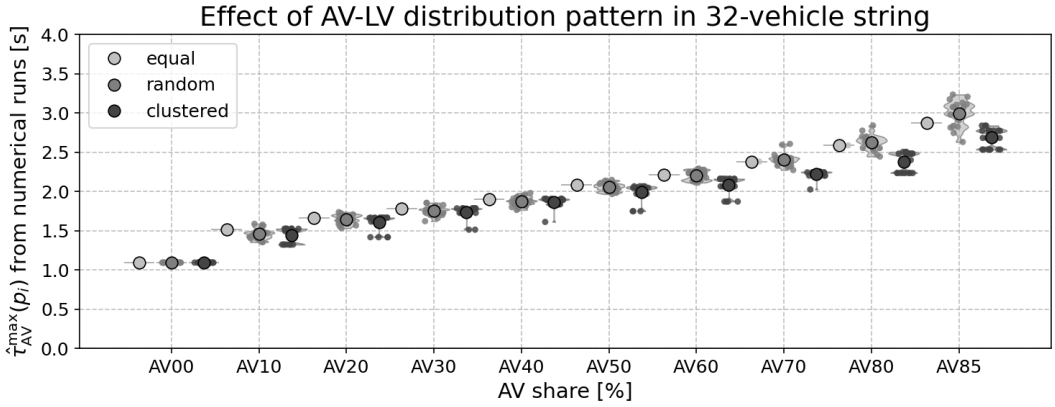

Figure 4.5: ToC-induced peak time headways, $\hat{\tau}_{\mathrm{AV}}^{\max}(p_i)$ by AV share for three within-platoon placement patterns in a single-lane string of 32 vehicles. **equal:** AVs distributed as evenly as possible along the string. **random:** ordering sampled uniformly at random. **clustered:** AVs arranged in contiguous blocks within the string. The violins summarise the distribution across permutations for each share and pattern.

capacity estimator to obtain traffic-mix-specific capacities, which were then compared with reductions measured in the multi-lane simulations. In the mid-range of AV shares, the estimator underpredicted the loss. The differences between placement patterns observed in Figure 4.5 are modest and cannot account for this gap. The principal source of the mismatch is therefore sought in mechanisms absent from the single-lane experiment, namely lane changes, lateral interactions and behavioural heterogeneity. Table 4.1 summarises the unmanaged mean time headways alongside the corresponding $\hat{\tau}_{\mathrm{AV}}^{\max}(p_i)$ values used for the capacity estimation.

	AV00	**AV10**	**AV20**	**AV30**	**AV40**	**AV50**	**AV60**	**AV70**	**AV80**	**AV85**
Unmanaged Means	1.58s	1.58s	1.90s	2.23s	2.46s	2.79s	3.03s	3.34s	3.51s	3.72s
Estimates $\hat{\tau}_{\mathrm{AV}}^{\max}(p_i)$	1.09s	1.51s	1.66s	1.78s	1.90s	2.08s	2.22s	2.38s	2.59s	2.88s

Table 4.1: Mean time headways [s] per traffic mix (unmanaged scenario) alongside $\hat{\tau}_{\mathrm{AV}}^{\max}(p_i)$ used in the capacity estimation of Paper III.

4.3.3 Linking ToR Distribution, Lane Changes, and Capacity

To further investigate the ToC-induced capacity reductions identified in Paper III, Figure 4.6 presents lane-change (LC) rates within the ODD zone, reported as differences relative to the no-ToC baseline and expressed per vehicle-hour at capacity. Panel (a) shows that the LC rate increases notably with AV share across traffic mixes.

a lesser degree relative to other factors.

Interestingly, at lower to mid shares (AV10–AV40) the managed case exhibits higher rates than the unmanaged case. From AV50 upwards, the unmanaged case exceeds the managed case and the gap widens with share. Panel (b) shows the LC rate along the ODD zone in distance bins for two exemplary traffic mixes (AV30–MV70 and AV70–MV30), confirming that LC activity is not uniform along the ODD zone. In managed runs, excess LCs are concentrated near the ODD start and decay along the zone with little activity near the ODD end. In unmanaged runs, excess LCs are more dispersed and show a build up towards the ODD end.

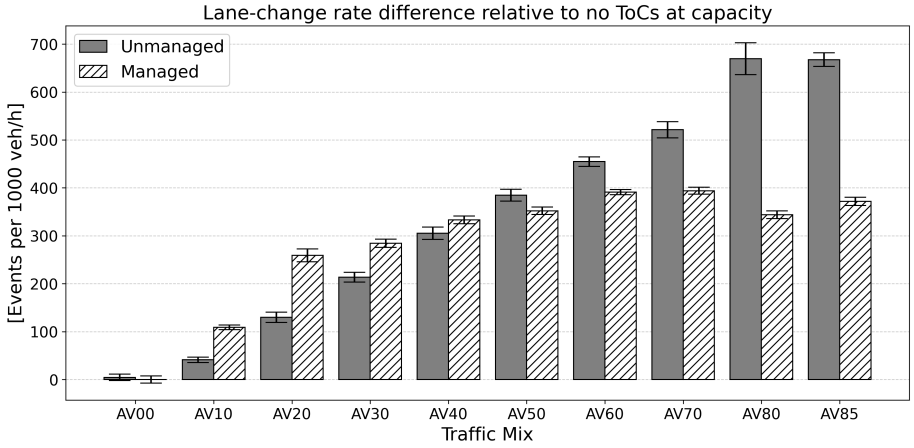

(a) Aggregated

(b) Binned

Figure 4.6: Lane-change rates within the ODD zone, normalised relative to the no-ToC baseline and expressed per vehicle-hour to account for throughput differences between Managed and Unmanaged. Rates are expressed as events per 1000 veh/h. **Panel (a)**: Difference in lane-change rate at capacity across traffic mixes. Bars show the mean across seeds; error bars indicate the standard deviation across seeds. **Panel (b)**: Lane-change rate along the ODD zone in bins for two exemplary traffic mixes (AV30–MV70 and AV70–MV30). Vertical lines indicate the ODD start and end.

As lane changes during the ToC preparation phase are not permitted for the respective individual AVs, the ToC-induced LC patterns are affected indirectly through reactive manoeuvres of surrounding vehicles. The mechanism within the ODD zone is posited as follows: (i) In the unmanaged case, the late concentration of ToCs near the ODD end is accompanied by a late concentration of reactive LCs in the capacity-critical end zone, producing local disruptions that propagate upstream and increase local speed and time headway variance. Consequently, a sustained downstream outflow deficit is detected near the ODD end, which aligns with the previously reported mid-range capacity reductions and the severe losses at higher shares. (ii) In the managed case, earlier ToR triggering shifts reactive LCs upstream, away from the ODD end, so even when excess LC counts are higher at low to mid shares, their timing and location render them less disruptive and, together with lower post-ToC MV time headways (per the AV/MV parameterisation), help to preserve higher throughput.

As the single-lane numerical experiments exclude lane changes, this lateral mechanism explains why the peak-headway estimator underpredicted capacity loss in the mid-range, and the observed rise in LC rate with AV share in unmanaged traffic is in line with the heavier right tails and higher means in the time headway distributions shown in Figure 4.4.

4.4 Summary of Findings and Limitations

4.4.1 Synopsis of Discussion Findings

For each of Papers I, II and III, complementary results were presented in the preceding sections to contextualise and qualify the original findings. These can be summarised as follows:

RL-based ToC management: Retrained policies yielded higher rewards relative to the baseline policy under the original formulation, yet their performance remained sensitive to reward balance and local minima. Additional PPO experiments clarified the trade-off between efficiency and throughput under varying weights. Overall, management of ToCs within an ODD zone appears preferable to unmanaged operation, and an optimal solution is suspected to be quantifiable only along a Pareto front.

Summary on RQ 1: Unmanaged ToCs markedly reduce efficiency in dense traffic, while ToR scheduling within the ODD mitigates the loss.

Safety under ODD constraints: Beyond the 60 km/h ODD limit use case reported in Paper II, complementary results were provided for the alternative No-AD zone use case to complete the safety assessment across the full AV share range. Per-hour SSM event rates showed the most severe ToC-induced impact in the mixed-autonomy range (AV20–AV50), whilst the elevated counts were dominated by inter-vehicle interactions. Intra-vehicle event rates were consistently lower. Management of ToCs within the ODD zone mitigated these detrimental safety impacts noticeably relative to unmanaged ToCs at the ODD end.

Summary on RQ 2: Safety impacts are conditional on the ODD boundary and the prevailing flow regime: in stable constant flow they peak at mid-range AV shares, driven by mixed-autonomy inter-vehicle interactions, and then decline as the AV share rises, whereas in accelerating flow they increase with AV share. ToR scheduling within the ODD reduces these effects in the stable-flow regime, while damping the time headway increase during ToCs reduces them in the accelerating-flow regime.

ToC-induced capacity losses: By comparing time headway distributions from the two-lane simulations in Paper III with synthesised headways from the single-lane numerical experiments, an underprediction of capacity reductions at intermediate AV shares was identified relative to measured losses. Further tests that varied intra-platoon AV–MV ordering yielded only minor differences in maximum time headways. Therefore, lane-change analyses at capacity were conducted. LC pattern differences between managed and unmanaged ODD zones clarify the interplay between ToC-induced flow disruption, upstream LC rates and capacity reductions, thereby accounting for the observed mismatch between the estimator and the simulation results.

Summary on RQ 3: ToCs constrain capacity within the ODD zone through increased effective time headways caused by preparation gaps and consequent lane-change-induced disruptions. Upstream ToR scheduling at the ODD boundary mitigates the constraint but does not restore baseline capacity.

4.4.2 Limitations, Scope and External Validity

The following points set out the principal constraints that qualify the interpretation of the results across the three main publications.

Scenario scope: The scenarios across the publications were designed to investigate different ODD boundary cases of Level 3 automated vehicles, particularly the ToC impact from automated to manual control. Nevertheless, it should not be assumed that all relevant usage contexts are covered. The results obtained for the 60 km/h ODD limit case are justified for Level 3 systems certified at that limit

under Regulation No. 157 (2021). At least one manufacturer has since upgraded to 95 km/h[15] following the 2023 revision of Regulation No. 157, which permits speeds up to 130 km/h. Papers I and III also analyse scenarios assuming operation up to 130 km/h (synthetic use case 5.1 in Paper I) and up to 100 km/h (calibrated motorway scenario in Paper III), although such high-speed systems are not expected by manufacturers before 2030. As Level 3 conditional driving is currently limited to motorway operation and is likely to remain so for years, geometry and the wider network context, for example urban operation, were not varied. Accordingly, the scope of inference is confined to these scenario settings and remains largely theoretical until supported by field evidence, so extrapolation to other contexts should be undertaken with caution.

Simulation tool fidelity: SUMO was used throughout, as it is presently the only microsimulation platform with a ToC model for network-level analysis. A known limitation is the representation of traffic stability: its car-following models are calibrated for collision avoidance and stable flow, which can damp instability phenomena of interest. The use of an ACC model as a proxy for Level 3 car-following behaviour cannot currently be verified against field evidence. In addition, SUMO lacks an explicit mechanism for human reaction time, which makes it difficult to calibrate to values often derived from driving-simulator studies. SUMO is also not designed to reproduce capacity drops on motorways with high accuracy[16]. Consequently, outcomes may overestimate safety and underestimate efficiency and capacity impacts.

Behavioural heterogeneity, compliance and ToC process assumptions: Given that all results depend on the AV share in the traffic mix, the usage rate in automated mode was assumed to be 100% and compliance with ToR was idealised. More importantly, the analyses were designed to isolate effects arising from automation during ToC preparation. Accordingly, human reaction times to ToRs were not varied. Instead, ToC durations were sampled from a skewed distribution with a mean of 7 s [17] (for non-emergency, system-initiated ToCs based on literature available at the time), although more recent studies suggest lower averages near 5 s[18]. In general, driver and AV controller heterogeneity was sampled from distributions across all influential car-following parameters. Direct calibration was not possible. Moreover, manufacturer procedures and field reference data for ToC timing, automation latencies and V2X availability were not publicly available, limiting external validation of

[15]The ODD of this system is currently bounded to the rightmost lane in a distinct car-following situation on motorways, without overtaking manoeuvres. This case was emulated in Paper III with complementary simulations under the tag "rightmost95".

[16]Paper III discusses the underlying mechanisms in more detail.

[17]By design, this parametrisation excludes lane-blockage effects from persistent MRMs. These were out of scope for this study.

[18]Paper III probed lower response times selectively and reported persistent ToC-related implications, yet no sensitivity analysis was undertaken to quantify PRT effects.

these assumptions.

Paper-specific methodological limits: In Paper I, the control task proved challenging to formalise in RL and resolve satisfactorily from a traffic-science perspective, as it required balancing traffic efficiency against the continuation of automated driving. This trade-off undermined training stability and, in turn, reduced the realised throughput in the scenario for many failed training runs. Consequently, overall performance remained sensitive to the reward balance, with RL training prone to local minima. Moreover, policy robustness was limited by the available training budget and by demand variability. In Paper II, the safety analysis was comparative and diagnostic rather than calibrated to absolute crash risk. Accordingly, fixed thresholds for the SSMs were not applied, as suitable calibrations are not available in SUMO. Specifically, MDRAC was chosen over the conventional DRAC for greater sensitivity. However, the PRT in the MDRAC denominator was held fixed at 1 s. Generally, as discussed extensively in the publication, inference from surrogate measures in simulation, rather than from direct criticality indicators such as collisions, is volatile and calibration dependent, which limits reliable quantification. Accordingly, the study reports relative changes in criticality. In Paper III, the peak headway estimator was constructed from single-lane numerical experiments to derive capacity estimates for each traffic mix. Those experiments exclude lane changes and lateral interactions, therefore the estimator is insensitive to lateral sorting and lane-change-induced perturbations that arise in multi-lane traffic. In the multi-lane simulations, unmanaged ToRs concentrated reactive lane changes near the ODD end and reduced throughput, whereas management moved this activity upstream. Because the estimator lacks this lateral mechanism, it underpredicted the capacity loss, particularly at intermediate AV shares. This mismatch was partly expected, but the extent of it, driven by the strong impact of lane-change manoeuvres, was underestimated. However, near 100% AVs, homogenisation of the flow reduces such lateral effects and the estimator aligns more closely with measured values, indicating that ToC-induced peak headways during control transitions contribute to capacity reductions, although quantifying the portion independent of lateral effects remains challenging.

Data, calibration and generalisability: As noted throughout, the principal limitation across all three studies is the absence of field reference datasets on Level 3 automated driving behaviour, including explicit ToC timing, together with undisclosed manufacturer procedures. Without such evidence, ToC model behaviour cannot be validated externally, and parameters relevant to the ToC process and automated system behaviour cannot be calibrated. V2X availability on the vehicle side, a prerequisite for managing ToRs via infrastructure elements such as RSUs, was not publicly disclosed, so the equipment rates assumed in the simulations are

idealised. Likewise, end-to-end communication latencies and reliability were not considered, so the practical feasibility of infrastructure-managed ToRs remains uncertain. Practical relevance also depends on market penetration and behavioural compliance. The effects quantified here arise at substantial AV shares under a 100% usage and ToR-compliance assumption, whereas such conditions are unlikely to be met in practice for some time. Over-the-air upgrades may expand ODD capabilities or alter ToC logic, which could mitigate or shift some of the reported effects, but without field evidence these pathways remain speculative. Accordingly, generalisability to early deployments of Level 3 systems is limited.

Chapter synthesis

This discussion synthesised complementary analyses and results from the three main studies, which yield a consistent picture: system-initiated, non-emergency ToCs impose measurable losses in dense traffic, and traffic management assistance can attenuate these impacts to a limited degree. The stated limitations delineate the envelope within which these claims hold. The following Chapter 5 distils these insights into concise conclusions and immediate implications for practice and research.

5 Conclusion

This thesis investigates how non-emergency, system-initiated ToCs in Level 3 automation, triggered by ODD limits, scale from individual vehicle behaviour to network-level outcomes. It utilises SUMO as the modelling and simulation framework, integrating an explicit ToC mechanism to examine efficiency, capacity, and safety across scenarios that progress from local dynamics to corridor- and network-level effects. The focus is not on the mere existence of ToCs, but on how their timing, location, and automation behaviour affect traffic once many such events arise under dense conditions. The studies form a coherent arc: the first explores ToR timing strategies and their immediate efficiency impacts; the second investigates safety-relevant effects in traffic flows at ODD boundaries; and the third quantifies capacity reductions and links preparatory ToC automation behaviour to those reductions. The central premise is that ToCs are recurring traffic events in conditional Level 3 automation.

In brief, three research questions ask to what extent and under which conditions ToCs affect efficiency, safety, and capacity, and whether targeted measures can alleviate the consequences. Regarding RQ 1, clustered ToCs in dense traffic disrupt flow disproportionately and, at high concentrations, can precipitate breakdown. Earlier work, Paper I, and complementary results show that preserving efficiency while sustaining automated operation within the ODD is a genuine trade-off rather than a free gain. TM that schedules and distributes ToCs, whether heuristic or RL-based, can mitigate delay by optimising along a Pareto front but it cannot restore baseline conditions. For RQ 2 on safety, the analyses indicate degraded safety performance at ODD boundaries in both quasi-stationary and accelerating traffic, particularly when ToCs cluster. While SSMs provide consistent directional evidence, precise magnitudes are constrained by SSM-based evaluation and threshold selection. TM attenuates these safety-relevant effects, with the largest gains at mid-range AV shares. Capacity, as posed in RQ 3, is reduced under all examined conditions in the presence of ToCs, and the reductions are more severe when ToCs cluster. Nevertheless, TM can partially mitigate them by minimising disruptions within a designated ODD corridor. Taken together, the answers support the overarching theme that non-emergency ToCs at ODD boundaries in conditional Level 3 driving are predictable yet unavoidable, and that design and V2X-supported TM can limit

their detrimental consequences but cannot make them fully recoverable.

In closing, one may reasonably ask what these findings might mean for current and future practice in transport. From a traffic perspective, many of the effects identified are likely to become materially relevant only at market penetration and usage rates that exceed what is realistic in the near term. For manufacturers, the principal challenge will remain on the automation side: ensuring sufficiently timely takeovers in increasingly complex and dynamic operational domains. How well and how quickly this can be achieved on the way to substantially broader market penetration of Level 3 can only be speculated upon here. A look at current developments in Level 4 automation, advanced in ever-larger testbed programmes worldwide yet still struggling with inadequate system reliability, leaves a rather sceptical impression. Level 3 functions with comparable operational domains will, under higher demands, have to demonstrate at least similar performance if they are to achieve sufficient attractiveness and acceptance among customers. From a human-centred perspective, the handover processes themselves are, at present, implemented in line with state-of-the-art research in human factors and current regulations with respect to comprehensibility, modality, driver readiness, and urgency. If anything, regulatory developments and adjustments that extend ODDs and enable better coordination through TM measures may spark regime shifts on both the automation and user sides and thereby alter the mechanisms documented here. The fundamental out-of-the-loop problem of automation will, however, not be fully resolved. Whether Level 3 systems, on the path towards fully autonomous driving, will establish themselves in the market and the traffic system in the long run, or instead amount to a brief staging post on the way to Level 4 and Level 5 with little discernible effect on traffic, is something only the future will show.

Bibliography

[1] Mohamed Abdel-Aty and Shengxuan Ding. "A matched case-control analysis of autonomous vs human-driven vehicle accidents". In: *Nature Communications* 15 (2024), p. 4931. DOI: 10.1038/s41467-024-48526-4.

[2] Robert Alms, Benjamin Couéraud, and Peter Wagner. "Perspectives on an ALKS Model in SUMO". In: *SUMO Conference Proceedings* 5 (July 2024), pp. 269–285. DOI: 10.52825/scp.v5i.1198.

[3] Robert Alms, Yun-Pang Flötteröd, Evangelos Mintsis, Sven Maerivoet, and Alejandro Correa. "Traffic Management for Connected and Automated Vehicles on Urban Corridors - Distributing Take-Over Requests and Assigning Safe Spots". In: *MFTS 2020 The 3rd Symposium on Management of Future Motorway and Urban Traffic Systems*. July 2020. URL: https://elib.dlr.de/134136/.

[4] Pablo Alvarez Lopez, Michael Behrisch, Laura Bieker-Walz, Jakob Erdmann, Yun-Pang Flötteröd, Robert Hilbrich, Leonhard Lücken, Johannes Rummel, Peter Wagner, and Evamarie Wießner. "Microscopic Traffic Simulation using SUMO". In: *The 21st IEEE International Conference on Intelligent Transportation Systems*. IEEE, Nov. 2018, pp. 2575–2582. URL: https://elib.dlr.de/127994/.

[5] *ASAM OpenODD Base Standard: Specification*. Standard 1.0.0. ASAM e.V., 2025. URL: https://publications.pages.asam.net/standards/ASAM_OpenODD/ASAM_OpenODD/latest/specification/index.html.

[6] James Ault and Guni Sharon. "Reinforcement Learning Benchmarks for Traffic Signal Control". In: *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*. Ed. by J. Vanschoren and S. Yeung. Vol. 1. 2021. URL: https://datasets-benchmarks-proceedings.neurips.cc/paper_files/paper/2021/file/f0935e4cd5920aa6c7c996a5ee53a70f-Paper-round1.pdf.

[7] Mike Blommer, Reates Curry, Radhakrishnan Swaminathan, Louis Tijerina, Walter Talamonti, and Dev Kochhar. "Driver brake vs. steer response to sudden forward collision scenario in manual and automated driving modes". In:

Transportation Research Part F: Traffic Psychology and Behaviour 45 (2017), pp. 93–101. ISSN: 1369-8478. DOI: /10.1016/j.trf.2016.11.006.

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

[9] California Department of Motor Vehicles. *Disengagement Reports*. Accessed 2025-09-04. 2025. URL: https://www.dmv.ca.gov/portal/vehicle-industry-services/autonomous-vehicles/disengagement-reports/.

[10] SC Calvert and Bart van Arem. "A generic multi-level framework for microscopic traffic simulation with automated vehicles in mixed traffic". In: *Transportation Research Part C: Emerging Technologies* 110 (2020), pp. 291–311. DOI: 10.1016/j.trc.2019.11.019.

[11] Simeon C. Calvert, Wouter J. Schakel, and J.W.C. van Lint. "A generic multi-scale framework for microscopic traffic simulation part II – Anticipation Reliance as compensation mechanism for potential task overload". In: *Transportation Research Part B: Methodological* 140 (2020), pp. 42–63. ISSN: 0191-2615. DOI: /10.1016/j.trb.2020.07.011.

[12] Tianshu Chu, Jie Wang, Lara Codecà, and Zhaojian Li. "Multi-Agent Deep Reinforcement Learning for Large-Scale Traffic Signal Control". In: *IEEE Transactions on Intelligent Transportation Systems* 21.3 (2020), pp. 1086–1095. DOI: 10.1109/TITS.2019.2901791.

[13] Alejandro Correa, Robert Alms, Javier Gozalvez, Miguel Sepulcre, Michele Rondinone, Robbin Blokpoel, Leonhard Lücken, and Gokulnath Thandavarayan. "Infrastructure Support for Cooperative Maneuvers in Connected and Automated Driving". In: *2019 IEEE Intelligent Vehicles Symposium (IV)*. 2019, pp. 20–25. DOI: 10.1109/IVS.2019.8814044.

[14] Hanwen Deng, Guoliang Xiang, Jiandong Pan, Xianhui Wu, Chaojie Fan, Kui Wang, and Yong Peng. "How to design driver takeover request in real-world scenarios: A systematic review". In: *Transportation Research Part F: Traffic Psychology and Behaviour* 104 (2024), pp. 411–432. ISSN: 1369-8478. DOI: /10.1016/j.trf.2024.06.012.

[15] Shengxuan Ding, Mohamed Abdel-Aty, Natalia Barbour, Dongdong Wang, Zijin Wang, and Ou Zheng. "Exploratory analysis of injury severity under different levels of driving automation (SAE Levels 2 and 4) using multi-source

data". In: *Accident Analysis & Prevention* 206 (2024), p. 107692. DOI: /10.1016/j.aap.2024.107692.

[16] Ebru Dogan, Mohamed-Cherif Rahal, Renaud Deborne, Patricia Delhomme, Andras Kemeny, and Jérôme Perrin. "Transition of control in a partially automated vehicle: Effects of anticipation and non-driving-related task involvement". In: *Transportation Research Part F: Traffic Psychology and Behaviour* 46 (2017), pp. 205–215. ISSN: 1369-8478. DOI: /10.1016/j.trf.2017.01.012.

[17] Hi-Drive. *Addressing challenges toward the deployment of higher automation (Hi-Drive)*. H2020 Research Project 101006664. European Commission, 2025. DOI: 10.3030/101006664. URL: https://cordis.europa.eu/project/id/101006664.

[18] Na Du, Feng Zhou, Elizabeth M. Pulver, Dawn M. Tilbury, Lionel P. Robert, Anuj K. Pradhan, and X. Jessie Yang. "Predicting driver takeover performance in conditionally automated driving". In: *Accident Analysis & Prevention* 148 (2020), p. 105748. ISSN: 0001-4575. DOI: /10.1016/j.aap.2020.105748.

[19] Na Du, Feng Zhou, Dawn M. Tilbury, Lionel P. Robert, and X. Jessie Yang. "Behavioral and physiological responses to takeovers in different scenarios during conditionally automated driving". In: *Transportation Research Part F: Traffic Psychology and Behaviour* 101 (2024), pp. 320–331. ISSN: 1369-8478. DOI: /10.1016/j.trf.2024.01.008.

[20] Yrvann Emzivat, Javier Ibanez-Guzman, Philippe Martinet, and Olivier H. Roux. "Dynamic Driving Task Fallback for an Automated Driving System whose Ability to Monitor the Driving Environment has been Compromised". In: *2017 IEEE Intelligent Vehicles Symposium (IV)*. 2017, pp. 1841–1847. DOI: 10.1109/IVS.2017.7995973.

[21] Alexander Eriksson and Neville A. Stanton. "Takeover Time in Highly Automated Vehicles: Noncritical Transitions to and From Manual Control". In: *Human Factors* 59.4 (2017), pp. 689–705. DOI: 10.1177/0018720816685832.

[22] ETSI. *Intelligent Transport Systems (ITS); Vehicular Communications; Manoeuvre Coordination Service (MCS); Pre-standardisation Study; Release 2*. Technical Report TR 103 578 V2.1.1. Accessed: 2025-07-30. European Telecommunications Standards Institute, 2024. URL: https://portal.etsi.org/webapp/WorkProgram/Report_WorkItem.asp?WKI_ID=53496.

[23] European Transport Safety Council (ETSC). *UN Working Party cautious on driving features that blur the lines between human and automated driving*. ETSC News. Accessed 2025-09-04. Jan. 2025. URL: https://etsc.eu/un-working-party-cautious-on-driving-features-that-blur-the-lines-between-human-and-automated-driving/.

[24] Scott Fujimoto, Herke van Hoof, and David Meger. "Addressing Function Approximation Error in Actor-Critic Methods". In: *CoRR* abs/1802.09477 (2018). arXiv: 1802.09477. URL: http://arxiv.org/abs/1802.09477.

[25] Christian Gold, Riender Happee, and Klaus Bengler. "Modeling take-over performance in level 3 conditionally automated vehicles". In: *Accident Analysis & Prevention* 116 (2018). Simulation of Traffic Safety in the Era of Advances in Technologies, pp. 3–13. ISSN: 0001-4575. DOI: /10.1016/j.aap.2017.11.009.

[26] Christian Gold, Moritz Körber, David Lechner, and Klaus Bengler. "Taking Over Control From Highly Automated Vehicles in Complex Traffic Situations: The Role of Traffic Density". In: *Human Factors* 58.4 (2016), pp. 642–652. DOI: 10.1177/0018720816634226.

[27] Timotej Gruden, Sašo Tomažič, and Grega Jakus. "Post-Takeover Proficiency in Conditionally Automated Driving: Understanding Stabilization Time with Driving and Physiological Signals". In: *Sensors* 24.10 (2024). ISSN: 1424-8220. DOI: 10.3390/s24103193.

[28] Magnus Gyllenhammar, Rolf Johansson, Fredrik Warg, Dejiu Chen, Hans-Martin Heyn, Martin Sanfridson, Jan Söderberg, Anders Thorsén, and Stig Ursing. "Towards an Operational Design Domain That Supports the Safety Argumentation of an Automated Driving System". In: *Proceedings of the 10th European Congress on Embedded Real Time Software and Systems (ERTS 2020)*. Toulouse, France, Jan. 2020.

[29] Ashley Hill, Antonin Raffin, Maximilian Ernestus, Adam Gleave, Anssi Kanervisto, Rene Traore, Prafulla Dhariwal, Christopher Hesse, Oleg Klimov, Alex Nichol, Matthias Plappert, Alec Radford, John Schulman, Szymon Sidor, and Yuhuai Wu. *Stable Baselines*. https://github.com/hill-a/stable-baselines. 2018.

[30] Lin Hu, Hai Cai, Jing Huang, Dongpu Cao, and Xin Zhang. "The Challenges of Driving Mode Switching in Automated Vehicles: A Review". In: *IEEE Transactions on Vehicular Technology* 73.2 (2024), pp. 1777–1791. DOI: 10.1109/TVT.2023.3319495.

[31] Mengxia Jin, Guangquan Lu, Facheng Chen, Xi Shi, Haitian Tan, and Junda Zhai. "Modeling takeover behavior in level 3 automated driving via a structural equation model: Considering the mediating role of trust". In: *Accident Analysis & Prevention* 157 (2021), p. 106156. ISSN: 0001-4575. DOI: /10.1016/j.aap.2021.106156.

[32] Thomas Jürgensohn. "Hybride Fahrermodelle". German. PhD thesis. Technische Universität Berlin, 1997.

[33] Zulqarnain H. Khattak, Michael D. Fontaine, and Brian L. Smith. "Exploratory Investigation of Disengagements and Crashes in Autonomous Vehicles Under Mixed Traffic: An Endogenous Switching Regime Framework". In: *IEEE Transactions on Intelligent Transportation Systems* 22.12 (2021), pp. 7485–7495. DOI: 10.1109/TITS.2020.3003527.

[34] Yan Kuang, Xiaobo Qu, Jinxian Weng, and Amir Etemad-Shahidi. "How Does the Driver's Perception Reaction Time Affect the Performances of Crash Surrogate Measures?" In: *PLOS ONE* 10.9 (Sept. 2015), pp. 1–13. DOI: 10.1371/journal.pone.0138617.

[35] Stefan Kupschick. "Modellierung menschähnlichen Fahrerverhaltens". German. Accepted version. Doctoral Thesis. Technische Universität Berlin, 2021. DOI: 10.14279/depositonce-11070.

[36] L3Pilot. *Piloting Automated Driving on European Roads (L3Pilot)*. H2020 Research Project 723051. European Commission, 2021. DOI: 10.3030/723051. URL: https://cordis.europa.eu/project/id/723051.

[37] Qingkun Li, Zhenyuan Wang, Wenjun Wang, Guofa Li, Jibo He, Liang Ma, and Bo Cheng. "Systematically modeling take-over performance: Considering the indirect effect of meteorological visibility mediated by drivers' attention". In: *Accident Analysis & Prevention* 220 (2025), p. 108174. ISSN: 0001-4575. DOI: /10.1016/j.aap.2025.108174.

[38] Zijian Lin and Feng Chen. "How various urgencies and visibilities influence drivers' takeover performance in critical car-following conditions? A driving simulation study". In: *Transportation Research Part F: Traffic Psychology and Behaviour* 104 (2024), pp. 303–317. ISSN: 1369-8478. DOI: /10.1016/j.trf.2024.06.007.

[39] Qingchao Liu, Jiaqi Liu, Yingfeng Cai, and Long Chen. "Exploring the Impact of the Takeover Time for Conditionally Automated Driving Vehicles on Traffic Flow in Highway Merging Area". In: *IEEE Transactions on Intelligent Transportation Systems* 23.12 (2022), pp. 24753–24764. DOI: 10.1109/TITS.2022.3202834.

[40] Zhenji Lu, Riender Happee, Christopher D.D. Cabrall, Miltos Kyriakidis, and Joost C.F. de Winter. "Human factors of transitions in automated driving: A general framework and literature survey". In: *Transportation Research Part F: Traffic Psychology and Behaviour* 43 (2016), pp. 183–198. ISSN: 1369-8478. DOI: /10.1016/j.trf.2016.10.007.

[41] Leonhard Lücken. *Resolving Collisions for the Gipps Car-Following Model*. 2019. arXiv: 1902.04927 [nlin.AO]. URL: https://arxiv.org/abs/1902.04927.

[42] Leonhard Lücken, Evangelos Mintsis, Kallirroi Porfyri, Robert Alms, Yun-Pang Flötteröd, and Dimitrios Koutras. "From Automated to Manual - Modeling Control Transitions with SUMO". In: *SUMO User Conference 2019*. Ed. by Melanie Weber, Laura Bieker-Walz, Robert Hilbrich, and Michael Behrisch. Vol. 62. EPiC Series in Computing. EasyChair, 2019, pp. 124–144. DOI: 10.29007/sfgk.

[43] Sven Maerivoet, Lars Akkermans, Kristof Carlier, Péter Pápics, Bart Ons, Stef Tourwé, Robert Alms, Yun-Pang Flötteröd, Leonhard Lücken, Evangelos Mintsis, Vasilios Karagounis, Dimitrios Koutras, Anton Wijbenga, Jaap Vreeswijk, Alejandro Correa, Xiaoyun Zhang, and Robbin Blokpoel. *TransAID Deliverable 4.2 – Preliminary Simulation and Assessment of Enhanced Traffic Management Measures*. Tech. rep. European Commission, Oct. 2020. DOI: 10.3030/723390. URL: https://cordis.europa.eu/project/id/723390/results.

[44] Gustav Markkula, Richard Romano, Ruth Madigan, Charles W. Fox, Oscar T. Giles, and Natasha Merat. "Models of Human Decision-Making as Tools for Estimating and Optimizing Impacts of Vehicle Automation". In: *Transportation Research Record* 2672.37 (2018), pp. 153–163. DOI: 10.1177/0361198118792131.

[45] Anthony D McDonald, Hananeh Alambeigi, Johan Engström, Gustav Markkula, Tobias Vogelpohl, Jarrett Dunne, and Norbert Yuma. "Toward computational simulations of behavior during automated driving takeovers: a review of the empirical and modeling literatures". In: *Human factors* 61.4 (2019), pp. 642–688. DOI: 10.1177/0018720819829572.

[46] Till Menzel, Gerrit Bagschik, Leon Isensee, Andre Schomburg, and Markus Maurer. "From Functional to Logical Scenarios: Detailing a Keyword-Based Scenario Description for Execution in a Simulation Environment". In: *2019 IEEE Intelligent Vehicles Symposium (IV)*. 2019, pp. 2383–2390. DOI: 10.1109/IVS.2019.8814099.

[47] Vicente Milanés and Steven E. Shladover. "Modeling cooperative and autonomous adaptive cruise control dynamic responses using experimental data". In: *Transportation Research Part C: Emerging Technologies* 48 (2014), pp. 285–300. ISSN: 0968-090X. DOI: /10.1016/j.trc.2014.09.001.

[48] Joel Andrew Miller, Soodeh Nikan, and Mohamed H. Zaki. "Navigating the Handover: Reviewing Takeover Requests in Level 3 Autonomous Vehicles". In: *IEEE Open Journal of Vehicular Technology* 5 (2024), pp. 1073–1087. DOI: 10.1109/OJVT.2024.3443630.

[49] Evangelos Mintsis, Dimitris Koutras, Kallirroi Porfyri, Evangelos Mitsakis, Leonhard Lücken, Jakob Erdmann, Yun-Pang Flötteröd, Robert Alms, Michele Rondinone, Sven Maerivoet, Kristof Carlier, Xiaoyun Zhang, Robbin Blokpoel, Martijn Harmenzon, and Steven Boerma. *TransAID Deliverable 3.1 - Modelling, simulation and assessment of vehicle automations and automated vehicles' driver behaviour in mixed traffic*. Tech. rep. European Commission, Sept. 2020. DOI: 10.3030/723390. URL: https://cordis.europa.eu/project/id/723390/results.

[50] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G. Bellemare, Alex Graves, Martin Riedmiller, Andreas K. Fidjeland, Georg Ostrovski, Stig Petersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran, Daan Wierstra, Shane Legg, and Demis Hassabis. "Human-level control through deep reinforcement learning". In: *Nature* 518.7540 (Feb. 2015), pp. 529–533. DOI: 10.1038/nature14236.

[51] Kristof Van Moffaert and Ann Nowé. "Multi-Objective Reinforcement Learning using Sets of Pareto Dominating Policies". In: *Journal of Machine Learning Research* 15.107 (2014), pp. 3663–3692. URL: http://jmlr.org/papers/v15/vanmoffaert14a.html.

[52] Walter Morales-Alvarez, Oscar Sipele, Régis Léberon, Hadj Hamma Tadjine, and Cristina Olaverri-Monreal. "Automated Driving: A Literature Review of the Take over Request in Conditional Automation". In: *Electronics* 9.12 (2020). ISSN: 2079-9292. DOI: 10.3390/electronics9122087.

[53] National Highway Traffic Safety Administration. *Standing General Order on Crash Reporting*. Originally issued 2021; accessed 2025-09-04. 2021. URL: https://www.nhtsa.gov/laws-regulations/standing-general-order-crash-reporting.

[54] National Highway Traffic Safety Administration. *AV TEST Initiative: Automated Vehicle Tracking Tool*. Accessed 2025-09-04. 2025. URL: https://avtest.nhtsa.dot.gov/.

[55] Frederik Naujoks, Christian Purucker, Katharina Wiedemann, and Claus Marberger. "Noncritical State Transitions During Conditionally Automated Driving on German Freeways: Effects of Non–Driving Related Tasks on Takeover Time and Takeover Quality". In: *Human Factors* 61.4 (2019), pp. 596–613. DOI: 10.1177/0018720818824002.

[56] OpenAI et al. "Dota 2 with Large Scale Deep Reinforcement Learning". In: *CoRR* abs/1912.06680 (2019). arXiv: 1912.06680. URL: https://arxiv.org/abs/1912.06680.

[57] Hengyan Pan, David B. Logan, Amanda N. Stephens, William Payre, Yonggang Wang, Zhipeng Peng, Yang Qin, and Sjaan Koppel. "Exploring the effect of driver drowsiness on takeover performance during automated driving: An updated literature review". In: *Accident Analysis & Prevention* 216 (2025), p. 108023. ISSN: 0001-4575. DOI: /10.1016/j.aap.2025.108023.

[58] Eleonora Papadimitriou, Omiros Athanasiadis, Gerdien Klunder, Simeon Calvert, Lin Xiao, and Bart van Arem. "A method to assess the safety implications of authority transitions in automated driving". In: *Traffic Safety Research* 6 (July 2024), e000048. DOI: 10.55329/fkix6369.

[59] Hannah Parr, Catherine Harvey, Gary Burnett, and Sarah Sharples. "Investigating levels of remote operation in high-level on-road autonomous vehicles using operator sequence diagrams". In: *Cognition, Technology & Work* 26.2 (2024), pp. 207–223. ISSN: 1435-5566. DOI: 10.1007/s10111-024-00762-w.

[60] Haorong Peng, Feng Chen, and Peiyan Chen. "Examining the Effects of Visibility and Time Headway on the Takeover Risk during Conditionally Automated Driving". In: *International Journal of Environmental Research and Public Health* 19.21 (2022). ISSN: 1660-4601. DOI: 10.3390/ijerph192113904.

[61] Sebastiaan M. Petermeijer, David A. Abbink, and Joost C. F. de Winter. "Should Drivers Be Operating Within an Automation-Free Bandwidth? Evaluating Haptic Steering Support Systems With Different Levels of Authority". In: *Human Factors* 57.1 (2015), pp. 5–20. DOI: 10.1177/0018720814563602.

[62] Linda Pipkorn, Marco Dozza, and Emma Tivesten. "Driver Visual Attention Before and After Take-Over Requests During Automated Driving on Public Roads". In: *Human Factors* 66.2 (2024), pp. 336–347. DOI: 10.1177/00187208221093863.

[63] Linda Pipkorn, Emma Tivesten, Carol Flannagan, and Marco Dozza. "Driver Response to Take-Over Requests in Real Traffic". In: *IEEE Transactions on Human-Machine Systems* 53.5 (2023), pp. 823–833. DOI: 10.1109/THMS.2023.3304003.

[64] Petr Pokorny and colleagues. "Descriptive analysis of reports on autonomous vehicle collisions in California: A 2021–mid-2022 update". In: *Traffic Safety Research* 1.2 (2022). URL: https://tsr.international/TSR/article/view/24451.

[65] Kallirroi N. Porfyri, Evangelos Mintsis, and Evangelos Mitsakis. "Assessment of ACC and CACC systems using SUMO". In: *SUMO 2018- Simulating Autonomous and Intermodal Transport Systems*. Ed. by Evamarie Wießner, Leonhard Lücken, Robert Hilbrich, Yun-Pang Flötteröd, Jakob Erdmann, Laura Bieker-Walz, and Michael Behrisch. Vol. 2. EPiC Series in Engineering. EasyChair, 2018, pp. 82–93. DOI: 10.29007/r343.

[66] Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. "Stable-Baselines3: Reliable Reinforcement Learning Implementations". In: *Journal of Machine Learning Research* 22.268 (2021), pp. 1–8. URL: http://jmlr.org/papers/v22/20-1364.html.

[67] Diederik M. Roijers, Peter Vamplew, Shimon Whiteson, and Richard Dazeley. "A survey of multi-objective sequential decision-making". In: *J. Artif. Int. Res.* 48.1 (Oct. 2013), pp. 67–113. ISSN: 1076-9757.

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

[69] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. "Proximal Policy Optimization Algorithms". In: *CoRR* abs/1707.06347 (2017). arXiv: 1707.06347. URL: http://arxiv.org/abs/1707.06347.

[70] Marios Sekadakis and George Yannis. "Systematic review and meta-analysis of take-over time from automated driving at SAE levels 2 and 3 to manual control". In: *Transportation Research Part F: Traffic Psychology and Behaviour* 113 (2025), pp. 263–306. ISSN: 1369-8478. DOI: /10.1016/j.trf.2025.04.003.

[71] Bobbie D. Seppelt and John D. Lee. "Modeling Driver Response to Imperfect Vehicle Control Automation". In: *Procedia Manufacturing* 3 (2015). 6th International Conference on Applied Human Factors and Ergonomics (AHFE 2015) and the Affiliated Conferences, AHFE 2015, pp. 2621–2628. ISSN: 2351-9789. DOI: /10.1016/j.promfg.2015.07.605.

[72] Elisabeth Shi and Klaus 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". In: *Accident Analysis & Prevention* 178 (2022), p. 106844. ISSN: 0001-4575. DOI: /10.1016/j.aap.2022.106844.

[73] Adam Skokan and Jan Mareček. "Could Disengagement Reports Indicate Evolution of Autonomous Vehicles?" In: *Vehicles* 7.2 (2025), p. 32. DOI: 10.3390/vehicles7020032.

[74] Xiaomei Tan and Yiqi Zhang. "A Computational Cognitive Model of Driver Response Time for Scheduled Freeway Exiting Takeovers in Conditionally Automated Vehicles". In: *Human Factors* 66.5 (2024), pp. 1583–1599. DOI: 10.1177/00187208221143028.

[75] Xiaomei Tan and Yiqi Zhang. "Understanding driver response to multi-stage takeover requests across varied modalities: A computational cognitive modeling approach". In: *Transportation Research Part C: Emerging Technologies* 174 (2025), p. 105114. ISSN: 0968-090X. DOI: /10.1016/j.trc.2025.105114.

[76] TransAID. *Transition Areas for Infrastructure-Assisted Driving (TransAID)*. H2020 Research Project 723390. European Commission, 2021. DOI: 10.3030/723390. URL: https://cordis.europa.eu/project/id/723390.

[77] Simon Ulbrich, Till Menzel, Andreas Reschka, Fabian Schuldt, and Markus Maurer. "Defining and Substantiating the Terms Scene, Situation, and Scenario for Automated Driving". In: *2015 IEEE 18th International Conference on Intelligent Transportation Systems*. 2015, pp. 982–988. DOI: 10.1109/ITSC.2015.164.

[78] UNECE. *Addendum 156 – UN Regulation No. 157 - Amendment 4 - Uniform provisions concerning the approval of vehicles with regard to Automated Lane Keeping Systems*. Tech. rep. Accessed on Mar. 12th, 2025. United Nations Economic commission for Europe, Mar. 2023. URL: https://unece.org/transport/documents/2023/03/standards/un-regulation-157-amend4.

[79] United Nations Economic commission for Europe (UNECE). *Addendum 156 – UN Regulation No. 157 - Uniform provisions concerning the approval of vehicles with regard to Automated Lane Keeping Systems*. url: https://unece.org/sites/default/files/2021-03/R157e.pdf Accessed on Nov. 28th, 2023. Mar. 2021. URL: https://unece.org/sites/default/files/2021-03/R157e.pdf.

[80] Rodolfo Valiente, Behrad Toghi, Ramtin Pedarsani, and Yaser P. Fallah. "Robustness and Adaptability of Reinforcement Learning-Based Cooperative Autonomous Driving in Mixed-Autonomy Traffic". In: *IEEE Open Journal of Intelligent Transportation Systems* 3 (2022), pp. 397–410. DOI: 10.1109/OJITS.2022.3172981.

[81] Silvia F. Varotto, Haneen Farah, Tomer Toledo, Bart van Arem, and Serge P. Hoogendoorn. "Modelling decisions of control transitions and target speed regulations in full-range Adaptive Cruise Control based on Risk Allostasis Theory". In: *Transportation Research Part B: Methodological* 117 (2018), pp. 318–341. ISSN: 0191-2615. DOI: /10.1016/j.trb.2018.09.007.

[82] Eugene Vinitsky, Aboudy Kreidieh, Luc Le Flem, Nishant Kheterpal, Kathy Jang, Cathy Wu, Fangyu Wu, Richard Liaw, Eric Liang, and Alexandre M. Bayen. "Benchmarks for reinforcement learning in mixed-autonomy traffic". In: *Proceedings of The 2nd Conference on Robot Learning*. Ed. by Aude Billard, Anca Dragan, Jan Peters, and Jun Morimoto. Vol. 87. Proceedings of Machine

Learning Research. PMLR, Oct. 2018, pp. 399–409. URL: https://proceedings.mlr.press/v87/vinitsky18a.html.

[83] Tobias Vogelpohl, Matthias Kühn, Thomas Hummel, and Mark Vollrath. "Asleep at the automated wheel—Sleepiness and fatigue during highly automated driving". In: *Accident Analysis & Prevention* 126 (2019). 10th International Conference on Managing Fatigue: Managing Fatigue to Improve Safety, Wellness, and Effectiveness"., pp. 70–84. ISSN: 0001-4575. DOI: /10.1016/j.aap.2018.03.013.

[84] Peter Wagner. "Analyzing fluctuations in car-following". In: *Transportation Research Part B: Methodological* 46.10 (2012), pp. 1384–1392. ISSN: 0191-2615. DOI: /10.1016/j.trb.2012.06.007.

[85] Changshuai Wang, Weilin Ren, Chengcheng Xu, Nan Zheng, Chang Peng, and Hao Tong. "Exploring the Impact of Conditionally Automated Driving Vehicles Transferring Control to Human Drivers on the Stability of Heterogeneous Traffic Flow". In: *IEEE Transactions on Intelligent Vehicles* 10.2 (2025), pp. 912–928. DOI: 10.1109/TIV.2024.3419789.

[86] Bradley W. Weaver and Patricia R. DeLucia. "A Systematic Review and Meta-Analysis of Takeover Performance During Conditionally Automated Driving". In: *Human Factors* 64.7 (2022), pp. 1227–1260. DOI: 10.1177/0018720820976476.

[87] Ran Wei, Anthony D. McDonald, Ranjana K. Mehta, and Alfredo Garcia. "Active Inference Models of AV Takeovers: Relating Model Parameters to Trust, Situation Awareness, and Fatigue". In: *Human Factors* 67.6 (2025), pp. 616–634. DOI: 10.1177/00187208241295932.

[88] Gereon Weiss, Marc Zeller, Hannes Schoenhaar, Christian Drabek Fraunhofer, and Andreas Kreutz. "Approach for Argumenting Safety on Basis of an Operational Design Domain". In: *2024 IEEE/ACM 3rd International Conference on AI Engineering – Software Engineering for AI (CAIN)*. 2024, pp. 184–193.

[89] Haoran Wu, Chaozhong Wu, Nengchao Lyu, and Jiannan Li. "Does a faster takeover necessarily mean it is better? A study on the influence of urgency and takeover-request lead time on takeover performance and safety". In: *Accident Analysis & Prevention* 171 (2022), p. 106647. ISSN: 0001-4575. DOI: /10.1016/j.aap.2022.106647.

[90] Wei Xue, Bo Yang, Tsutomu Kaizuka, and Kimihiko Nakano. "A Fallback Approach for an Automated Vehicle Encountering Sensor Failure in Monitoring Environment". In: *2018 IEEE Intelligent Vehicles Symposium (IV)*. 2018, pp. 1807–1812. DOI: 10.1109/IVS.2018.8500392.

[91] Chao Yu, Akash Velu, Eugene Vinitsky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and YI WU. "The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games". In: *Advances in Neural Information Processing Systems*. Ed. by S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh. Vol. 35. Curran Associates, Inc., 2022, pp. 24611–24624. URL: https://proceedings.neurips.cc/paper_files/paper/2022/file/9c1535a02f0ce079433344e14d910597-Paper-Datasets_and_Benchmarks.pdf.

[92] Jing Yu and Feng Luo. "Fallback Strategy for Level 4+ Automated Driving System". In: *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. 2019, pp. 156–162. DOI: 10.1109/ITSC.2019.8917404.

[93] Bo Zhang, Joost de Winter, Silvia Varotto, Riender Happee, and Marieke Martens. "Determinants of take-over time from automated driving: A meta-analysis of 129 studies". In: *Transportation Research Part F: Traffic Psychology and Behaviour* 64 (2019), pp. 285–307. ISSN: 1369-8478. DOI: /10.1016/j.trf.2019.04.020.

[94] Lai Zheng, Tarek Sayed, and Mohamed Essa. "Validating the bivariate extreme value modeling approach for road safety estimation with different traffic conflict indicators". In: *Accident Analysis & Prevention* 123 (2019), pp. 314–323. ISSN: 0001-4575. DOI: /10.1016/j.aap.2018.12.007.