A Data-Driven Integrated Framework for Virtual Testing of Autonomous Vehicles in Mixed Traffic Scenarios

Brunella Caroleo(✉), Javad Sadeghi, Cristiana Botta, Shadi Nikneshan, and Maurizio Arnone

Fondazione LINKS, Via P.C. Boggio 61, 10138 Turin, Italy
{brunella.caroleo,javad.sadeghi,cristiana.botta,shadi.nikneshan, maurizio.arnone}@linksfoundation.com

Abstract. The increasing presence of autonomous vehicles (AVs) in urban environments introduces both opportunities and challenges, particularly regarding their interactions with traditional vehicles and other road users. This paper presents a comprehensive framework designed to assess the integration of AVs in mixed traffic scenarios. The framework is built upon real-world data collected from AV trials conducted in Turin, Italy. By leveraging traffic microsimulation along with machine learning techniques, the study proposes a framework aimed at assessing ex-ante the impacts on traffic of AVs introduction, thus constituting a relevant tool of virtual testing of CCAM (Cooperative, Connected, and Automated Mobility) trials before the physical introduction of autonomous vehicles on public roads. The integration of High-Performance Computing (HPC) ensures the efficiency of these simulations, enabling real-time analysis and testing. The proposed framework not only provides decision-makers with a tool for virtual testing of AV deployment, but also offers actionable insights into traffic management strategies. The study's findings contribute to a deeper understanding of the role AVs can play in future urban mobility systems, particularly as cities prepare for the broader adoption of CCAM technologies.

Keywords: Cooperative Connected and Automated Mobility (CCAM) · Autonomous Vehicles (AV) · Traffic Management · Traffic simulation · Virtual Testing · Urban Transportation · Machine Learning

1 Introduction

Urban mobility is experiencing an era of transformation with the rise of Autonomous Vehicles (AVs) and Cooperative, Connected, and Automated Mobility (CCAM). While early research focused on private AVs, recent studies have shifted toward public and shared autonomous transportation, particularly autonomous shuttles using minibus technology. These shuttles are expected to become an essential element of future urban mobility, integrated within public transport systems.

A. Kocian et al. (Eds.): INTSYS 2024, LNICST 608, pp. 293–310, 2025.
https://doi.org/10.1007/978-3-031-86370-7_18

Current AV experiences, especially with shuttles, are often limited to controlled environments. The main challenge is to integrate these shuttles into mixed traffic, where they interact with pedestrians, cyclists, and traditional human-driven vehicles. Limited experiences with autonomous shuttles in mixed traffic conditions can be exemplified by initiatives such as the demonstrations within the European project SHOW (SHared automation Operating models for Worldwide adoption) [1], funded under the H2020 program. Transitioning from human-driven to autonomous vehicles requires careful planning to ensure safety, efficiency, and public acceptance. During this shuft, where autonomous trials are often geographically constrained (e.g. authorized routes), it is vital to evaluate their impact on urban traffic. However, the risk and feasibility analyses typically conducted by vehicle providers only partially address this need.

The city of Turin (Italy) serves as a vibrant environment for autonomous mobility studies, boasting a local CCAM ecosystem that has actively evolved in recent years. This ecosystem is a result of collaborative efforts, drawing on research funding from prominent initiatives such as the European SHOW project [1], the European IN2CCAM project [2], and the national ToMove project [3]. The city acts as a living lab, yielding insights into AV behaviour across various urban settings.

This work assumes a dual objective. Firstly, to provide a comprehensive framework that leverages actual data from urban infrastructure and autonomous vehicle trials, realistically simulating the interactions of these vehicles with other road users and infrastructure. Secondly, to provide decision-makers with a tool for virtual testing of AV trials before their physical deployment, thus preventing errors that may jeopardize urban mobility safety and assessing the impact of introducing AVs in specific areas before initiating live trials.

2 Related Works

The exploration of Autonomous Vehicles (AVs) and their integration into urban mobility is a multifaceted endeavour that has captivated the attention of researchers, engineers, and urban planners worldwide. Considerable research has been dedicated to understanding the implications of AVs in urban environments. Early studies primarily focused on the technological aspects, such as sensor fusion, perception algorithms, and decision-making processes of autonomous systems [4]. As AV technology matured, attention shifted towards assessing the impact of these vehicles on traffic flow, congestion, and overall urban mobility [5–7].

The specific domain of autonomous shuttle services, especially in public transportation, has garnered interest in recent years. Initiatives such as the European SHOW project [1] have explored the integration of autonomous shuttles into public transport systems [8–10], emphasizing their role in enhancing last-mile connectivity [11].

Traffic simulation has been used in numerous simulation research scenarios of AVs, highlighting its versatility and wide-ranging applications: it provides a helpful tool to answer complex research questions, to evaluate or test traffic management strategies and their impacts [12]. Simulation offers a safe and cost-effective method of testing and validating autonomous algorithms and control systems before real-world deployment. These simulations provide a crucial bridge between theory and practice, allowing researchers to refine AVs technologies in controlled virtual environments. Traffic simulators replicate situations and evaluate the consequences of different strategies for using AVs. In the majority of microscopic simulation research, AVs have been distinguished from human-driven vehicles based on their driving patterns, highlighting the pivotal role of driving behaviour in modelling AVs [13]. Understanding the AV driving behaviour is essential for creating realistic and effective simulations that mirror real-world conditions. Using microsimulation, some researchers studied the impact of different AVs penetration on traffic parameters [14] and on urban traffic flow and road capacity [7], highlighting potential risk that if AVs increase overall car use, they could strain traffic management. Additionally, microsimulation has emerged as a pivotal tool for understanding the complex interactions between AVs and other road users in realistic traffic scenarios. Existing works delve into the development of microsimulation environments for AVs, often utilizing platforms such as SUMO (Simulation of Urban MObility) [15]. The assessment of impacts related to autonomous services, safety, and overall traffic dynamics becomes feasible through this simulation, guiding the identification of optimal strategies in various contexts.

The safety of AVs and their interactions with pedestrians and traditional vehicles have been extensively investigated. Studies have utilized surrogate safety measures like Time to Collision (TTC) and Post Encroachment Time (PET) to evaluate and compare the safety of AV-pedestrian interactions with traditional vehicle-pedestrian interactions [16, 17].

Research efforts have also been made to integrate different simulation platforms to create comprehensive testing environments for AVs. The co-simulation of CARLA [18] (for the ego-AV simulation) and SUMO (for surrounding traffic microsimulation) has gained attention, allowing for more realistic testing scenarios [19, 20].

In summary, the landscape of AV research is vast and continually evolving, but research contributions in the field of virtual testing for autonomous vehicles have been sectoral, and an overall view is lacking. There are some scientific contributions that try to give an overview by cross-referencing the contributions that may come from different domains [21], but there is a lack of real data from connected/autonomous vehicles, and thus no concrete results from such integrated frameworks. In this paper, the authors propose to cover this gap by providing a comprehensive framework and testing it on real data collected in the field during autonomous driving experiments in Turin in recent years.

3 The Case Study

The City of Turin (Italy) serves as a dynamic setting for this study, with a robust CCAM ecosystem encompassing local institutions, research centres, the Traffic Control Center, and the public transport operator. In late 2022, as part of the H2020 SHOW project [1], two autonomous shuttles were tested in mixed traffic in Turin, making it the only Italian site among 20 European cities involved. These trials aim to innovate urban mobility and gather valuable data on autonomous vehicle behaviour in real-world traffic, guiding future autonomous experiments in the city [2, 3].

The Turin pilot was designed to offer a flexible, on-demand transport service with two SAE level 4 autonomous shuttles (provided by Navya [22]), which circulated on an authorized 5 km route near the City of Health and Science of Turin. The study's data analysed in this paper was collected during the pre-demonstration phase (August–October 2022), focusing on technical validation before the shuttles began carrying passengers. Navya's electric shuttle, capable of holding 15 passengers, uses LIDAR and localization technologies to navigate autonomously at speeds up to 18 km/h. To ensure safety, it operates in environments where traffic does not exceed 50 km/h. The shuttle slows or stops when an obstacle enters its "priority zone," typically positioned 1 m ahead of a checkpoint, and resumes only when the path clears.

During the test period, some videos were recorded across different days to observe interactions between pedestrians and autonomous shuttles at an unsignalized crosswalk with moderate traffic. This setup, free of traffic signal constraints (Fig. 1), allowed the study of natural interactions between road users, including pedestrians, conventional vehicles, buses, and the autonomous shuttle.

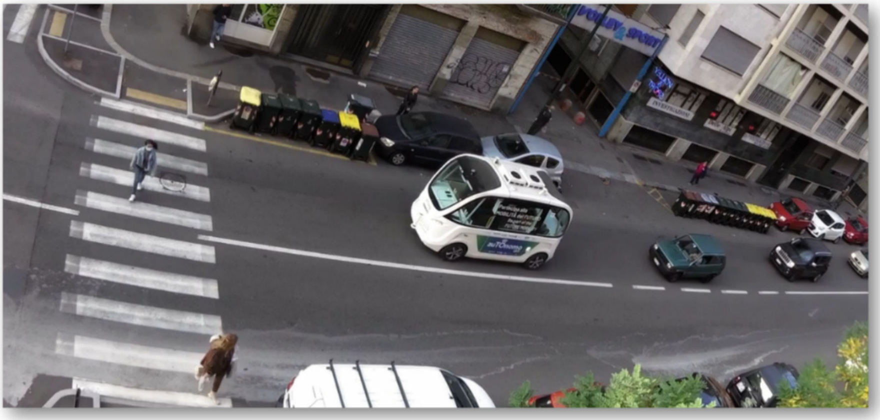

Fig. 1. Screenshot of the area chosen for video recording during the trial.

4 Methodology

The proposed framework is designed for adaptable use across various urban settings, adjusting to available data and traffic conditions, making it applicable beyond the specific case of Turin.

It consists of three main phases: data gathering, data analysis and modeling, and traffic microsimulation (Fig. 2). This modular structure allows each phase to function independently or integrate with others, providing flexibility across diverse urban contexts. In the Data Gathering phase, the framework adapts to a wide range of data sources, including AV sensors and infrastructure inputs. The Data Analysis and Modeling phase processes and validates this information, enabling machine learning applications to extract insights into AV behavior and interactions. The Traffic Microsimulation phase, using customizable models, simulates urban scenarios and can be implemented on various platforms, such as SUMO, tailored to project needs.

In the case-specific implementation for Turin case study, the framework incorporated local AV trial data, including sensor and traffic signal information, to create realistic microsimulations. This tailored approach accommodates Turin's unique urban characteristics, such as complex road networks and diverse traffic patterns.

By customizing each phase to align with specific data and technical resources, the framework offers a robust tool for assessing AV deployment impacts, applicable not only to Turin but also to other cities with similar data availability and mobility needs.

Fig. 2. The proposed framework.

4.1 Data Gathering

The foundation of the proposed framework lies in the extensive and detailed data collection process. Typically, the data sources of CCAM experimentations are:

1. Data from the autonomous vehicle itself, e.g. real-time location data, data from the IMU (Inertial Measurement Unit), current speed and acceleration, data of vehicle's steering and braking systems.
2. Data from the infrastructure, e.g. real-time traffic light phases (i.e. time to green, time to red).
3. Road network data, e.g. road network geometry, traffic demand, traffic signals, public transport routes and stops.
4. Sometimes, external camera captures (e.g. on the road) are available, and this data could be used to monitor behaviours of road users (i.e. traditional vehicles, autonomous vehicles and VRUs such as pedestrians and bikers) at critical road segments (e.g. pedestrian crossings). Keep in mind that vehicle providers often do not make available the videos taken by the cameras in the vehicle, so such information on interactions is not available unless external cameras filming the scenes are available.

In Turin case study presented in this paper the available data was the following:

1. Data from the autonomous vehicle itself: the Italian Smart Road Decree (regulating experimentations involving autonomous vehicles in Italy) asks at least for the following categories of data for autonomous experimentations taking place in Italy: date,

time, position in WGS84 coordinates; current operation mode (autonomous or manual); instant speed; instant acceleration; distance travelled since the beginning of the experimentation; activation of controls for the lateral dynamics of the vehicle; activation of controls for the longitudinal dynamics of the vehicle; number of revolutions per minute of the engine; gear ratio engaged, or other equivalent indicator; current value of yaw angle, roll and pitch; use of lighting and visual and acoustic signalling devices; acquired data of the sensors that are part of the system being tested; any V2V and V2I messages received and transmitted.

For each of the abovementioned category, data of the two autonomous shuttles has been collected for a period of three months (August–October 2022) for each second of the experimentation (data collected at a frequency of 1Hz).

2. Data from the infrastructure: all the traffic lights of the experimentation area (about 14 groups of traffic lights) were connected to the Traffic Control Centre of the City, and sent to the shuttle real-time traffic light phases (i.e. time to green, time to red) via SPATEM (Signal Phase And Timing Extended Message).
3. Road network data: the Municipality of Turin, in collaboration with the Traffic Operation Centre, put at disposal for the authorized zone: the road network geometry (crossed with Open Street Map), the traffic demand, information about traffic signals, and - thanks to the open data of the local Public Transport Operator – also public transport routes and stops were available via GTFS format.
4. An external camera has been installed at an unsignalized crosswalk and has recorded, in different days, about 10 h of interactions among different road users in the scene: traditional vehicles, autonomous vehicles, bikers and pedestrians, the last ones being free to cross that intersection without any traffic light regulation. The traffic scene was recorded by alternating two Garmin VIRB ™ Action Cams, with a resolution of 1080p HD and 30 frame/s. They were placed on the top of a carbon fibres telescopic pole at 10.80 m to the ground. The cameras were placed outside the roadway in a position that was difficult for drivers to detect.

4.2 Data Analysis and Modelling

Once data is collected in the previous step, the Data analysis and modelling phase includes the following activities:

- Data validation, pre-processing and exploration: this can be done by using statistical and machine learning techniques depending on the objectives of the data exploration. Included in this phase, if needed, compliance with legislative requirements is checked.
- Data visualization, to ease the interpretation of AVs data. This can be done -according to the objectives of the analysis- in real time (e.g. visualizing data of the autonomous vehicle collected via an API connection) or offline.
- Computation of distributions of parameters for next steps of analysis (e.g. speed profile for microsimulation setting and validation by vehicle type).
- If videos are available (source number 4 of the data in Sect. 4.1), data enrichment for a better comprehension of AVs behaviour with respect to the surrounding environment. In fact, in case of AVs circulating in mixed traffic and not reserved in dedicated lanes, it would be relevant to study AV behaviour with respect to other road users. This can be pursued by means of detection and tracking algorithms aimed to study the behaviours

of AVs and other road users in real traffic, and jointly with safety measures analysis of these interactions.

In Turin case study presented in this paper, validation of high-frequency monitoring data from AV sensors has been undertaken, ensuring the reliability of the dataset. Anomaly detection algorithms are implemented to uncover unexpected events, adding a layer of robustness to the dataset. The application of machine learning extends to unravelling hidden relationships among various features, providing valuable insights into the intricate dynamics of AV operations. Furthermore, compliance with legislative requirements has been systematically checked to ensure adherence to safety and operational standards.

A web dashboard in MS PowerBI has been created to visualize AV data, offering a user-friendly, interactive interface for both real-time (via API) and offline analysis (using.csv files from the vehicle provider). This setup allows analysts to check and interpret collected data. In this framework, data analysis (in Python) and data visualization/exploration (in PowerBI) serve distinct roles. Data analysis is used independently in Python to preprocess, validate, and analyze AV data, detecting anomalies and behavior patterns, such as outliers, speed changes, and potential traffic conflicts. The PowerBI dashboard, meanwhile, visualizes key metrics and trends, enabling stakeholders to explore data insights. Although data analysis results are not integrated in real-time, they are periodically uploaded to PowerBI as structured datasets (e.g., CSVs) for easy viewing. This setup ensures PowerBI is an accessible tool for insights derived from machine learning models, adaptable to both live and historical data analysis.

Given the absence of dedicated lanes for AVs in the case study area, understanding their behaviour in interaction with other road users is imperative. The methodology employs data enrichment by integrating AV sensors' data with information from cameras along the route. Due to the camera wide angle, the raw videos were affected by distortion error. The correction was performed according to the specific distortion matrix of the camera. Only the video sections including an interaction between pedestrian and autonomous/traditional vehicles were considered for further analysis. Video analysis, coupled with detection and tracking algorithms, offers a nuanced understanding of AV behaviour concerning the surrounding environment. Videos were collected across different days, and road users in the scene (i.e., pedestrians and vehicles) have been detected and tracked (Fig. 3). The object detection was performed using a customized YOLOv7 model [23] integrated with SORT algorithm [24]. This integration assigns each detected object a unique identifier that remains consistent across all frames, enabling seamless tracking of both pedestrians and autonomous/traditional vehicles in the video sequences.

Then, an algorithm automatically derived conflict measures to assess safety in conflicts, based on their spatial-temporal trajectories. The resulting distributions were compared through Kolmogorov-Smirnov (K-S) test for both traditional and autonomous vehicles. Notably, a specific emphasis is placed on studying surrogate safety measures (TTC-Time to Collision, PET-Post Encroachment Time) in AV-pedestrian interactions versus traditional vehicle-pedestrian interactions, revealing insights into the safety aspects of AV systems. Data distributions were modelled to understand differences regarding pedestrians' behaviours in front of the two vehicle types. The computation of distributions

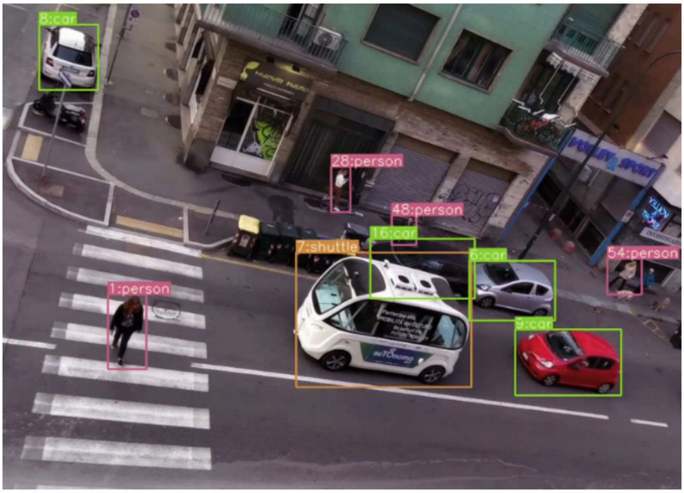

Fig. 3. Example of the detection and tracking algorithm applied to the video.

of parameters for microsimulation setting and validation by vehicle type, particularly focusing on speed profiles, contributes to a deeper comprehension of road users' patterns.

4.3 Traffic Microsimulation

After data collection and analysis, the framework includes a traffic microsimulation phase, aimed to assess to what extent traffic in a zone selected for AV testing is impacted by the introduction of autonomous vehicles on public roads. Traffic microsimulation, rather than macro or meso-simulation, is needed to simulate individual vehicle movements and interactions. Commonly used microscopic simulators include VISSIM, AIMSUN, and SUMO, which can be customized for specific scenarios to assess various AV deployment strategies. This virtual testing environment enhances efficiency by reducing the need for extensive road tests, providing quantitative insights into AVs' traffic impacts. While traffic microsimulation can assess AV effects on traffic, it lacks the ability to model AV-specific features, like sensor positions. To address this, integrating traffic simulators with autonomous driving simulators like CARLA (Car Learning to Act) [22] through co-simulation creates a more complete testing environment, aiding in AV development, testing, and validation. This combined approach helps AVs handle complex real traffic scenarios safely and reliably, while also providing realistic traffic edge scenarios based on observed data. Such simulations accelerate AV technology's readiness for deployment on public roads.

In the Turin case study the open-source simulator SUMO (Simulation of Urban Mobility) [15] was selected for traffic microsimulation, partly due to its active user and developer community, which supports continuous advancements and knowledge sharing. Using TraCI (Traffic Control Interface) with SUMO allowed real-time control and customization of traffic scenarios, bridging Python with SUMO for enhanced flexibility. This approach enables a wide range of applications, from testing traffic management strategies to developing autonomous vehicle algorithms, making SUMO particularly effective compared to other simulators.

To integrate traditional and autonomous vehicles in simulation, the key differences lie in decision-making and sensory inputs. While humans rely on visual cues, which can be subjective, AVs use sensors like LiDAR, radar, and cameras to gather objective data and respond instantaneously. These differences influence AV simulation parameters, such as shorter reaction times and reduced minimum gaps, reflecting the faster response capabilities of AVs compared to human drivers. This distinction allows more accurate modeling of AV behaviors in mixed traffic environments.

For traffic simulation of the autonomous vehicles involved in Turin case study, the network of the testing area has been imported from Open Street Map by OSM Web Wizard, and then some adjustments have been performed to be in accordance with the real road network (e.g. the number of lanes, traffic signal cycles, speed limits). The autonomous shuttle, with its own characteristics (e.g. length, speed profile, etc.), has been modelled in SUMO, thanks to the analysis performed in the previous step (Data analysis and modelling) as well as SUMO's customizability. In order to test different traffic scenarios, traffic and pedestrian flows, public transport routes and stops have been added to the simulation environment. Additionally, different traffic management and control strategies have been tested to assess their impacts. In implementing traffic scenarios in SUMO, after creating the simulation network and traffic flows, the TraCI was used to define scenarios and their associated logic. TraCI continuously monitors simulation to execute these logics, specifying when an intervention starts, what will happen during interventions, where it is located within the network, and when it will end. These logics control characteristics of network elements, such as vehicle permissions on specific lanes or streets, intersection control management, and vehicle behaviours and routes. This simulation framework prepared a powerful environment to define complex and dynamic scenarios by TraCI. This enables detailed analysis and testing of various scenarios and analysis without extensive real-world testing, which can be crucial for further development and optimization of autonomous vehicles' technology.

High-Performance Computing (HPC) infrastructure was used to meet the computational demands of real-time analytics and high-traffic networks. HPC allows rapid processing of large simulations, completing in minutes what might take hours or days on standard computers.

The framework also implemented a SUMO-CARLA co-simulation to integrate realistic traffic flows into autonomous driving simulations. Running on HPC machines, this co-simulation combines SUMO's urban traffic modeling with CARLA's detailed AV simulations, creating a comprehensive platform for AV research. SUMO manages traffic flow, transforming background vehicles for CARLA, while CARLA's AVs interact with traffic to complete driving tasks. By distributing tasks between CARLA and SUMO, the framework evaluates AV algorithms and protocols on both individual and traffic levels. This integration enables realistic testing of AV control, perception, and localization in a simulated environment before actual road trials. It enhances AV safety and reduces development costs by completing most validations within a controlled virtual space.

The resulting integrated framework has several purposes. Firstly, it enables virtual testing of AV deployment in real traffic conditions, fostering safe and controlled experiments. The development process becomes cost-effective, and the deployment of AV demonstrations is accelerated. Furthermore, the integration enhances scalability and transferability, ensuring that insights gained from the simulation environment can be effectively applied to diverse urban contexts. In conclusion, the employed methodology reflects a holistic and meticulous approach, incorporating cutting-edge technologies and methodologies to address the challenges posed by the integration of CCAM technologies into urban mobility scenarios.

5 Results

The implementation of the proposed framework has yielded insightful results across various dimensions, ranging from the validation of collected data to the impacts of AV circulation on urban traffic. This section delves into the key outcomes derived from each phase of the methodology, shedding light on the efficacy and relevance of the researchers' approach.

The first result, coming from the data gathering phase, is the data itself: the extensive data collection process spanning three months has proven to be a valuable asset for the proposed work. The dataset, encompassing diverse parameters related to AV operations and the urban road network, offers a comprehensive snapshot of real-world scenarios. Ongoing projects (IN2CCAM [2], ToMove [3]) are poised to contribute additional CCAM data from real experimentations, ensuring a continuous and evolving dataset for future research endeavours. This wealth of data serves as a foundation for subsequent phases of the methodology, enhancing the robustness and applicability of the framework.

As regards the data analysis and modelling phase, a relevant result comes from the data visualization. The deployment of a web dashboard for data visualization (Fig. 4) enhances the interpretability of AV data, fostering a user-friendly interface for stakeholders and researchers. The primary purpose of the PowerBI web dashboard is to facilitate comprehensive data analysis and visualization, enhancing understanding and decision-making related to shuttle operations. Additionally, the dashboard provides real-time data monitoring capabilities with flexible filtering options such as date, operating zone, and shuttle identification. It has been designed with four key pages to provide insights into vehicle performance, sensor data, signaling devices, and overall fleet operations:

(1) Vehicle Dynamics: this section displays crucial metrics such as speed, steering rate and angle, bearing, and battery level. These parameters are analysed over time to understand the shuttle's dynamic behaviour. Additionally, the percentage of time spent in various operational modes is provided, offering insights into how often the vehicle operates in autonomous mode.
(2) Sensors: on this page, various sensor readings are visualized, including GNSS correction status, door status, vehicle mode, and battery status. Temperature data (engine, indoor, outdoor) is also displayed. This page helps in monitoring and assessing the shuttle's operational conditions and overall performance.
(3) Lighting, Visual, and Acoustic Signalling Devices: this section focuses on the shuttle's signalling systems, such as the state of blinkers, brake lights, and reverse lights.

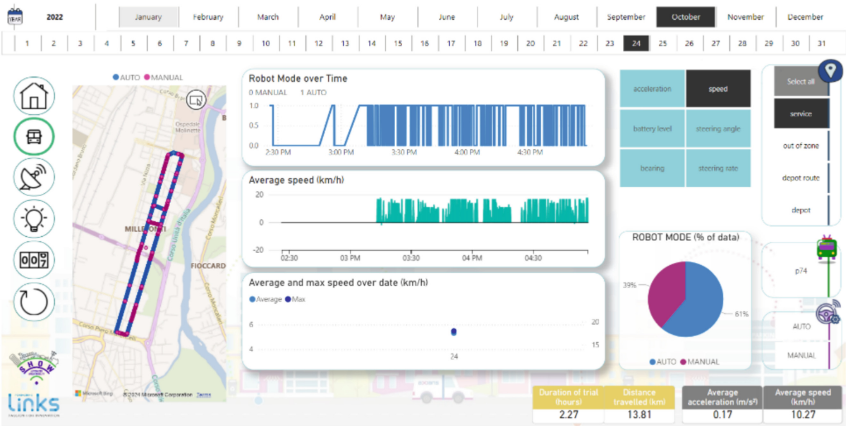

Fig. 4. Visualization and analysis of AVs data in PowerBI.

The analysis here aims to ensure that these critical safety components are functioning correctly, providing a statistical representation of their operation.

(4) Fleet Overview: designed for high-level insights, this page aggregates data across all shuttles in the fleet, showing key metrics like travel distance and duration for both manual and autonomous modes. It allows for comparisons across different shuttle IDs and operational modes, helping stakeholders to evaluate performance and compliance with legislative requirements.

Furthermore, the application of machine learning techniques in the analysis and modelling phase has yielded promising results, with anomaly detection activities mainly focused on detecting unusual behaviours in the data collected from autonomous vehicles' sensors. The main emphasis was placed on identifying abrupt speed changes (both acceleration and deceleration), which can indicate anomalies in vehicle operation. These anomalies are critical for ensuring the safety and efficiency of AVs in urban traffic. The primary method used for detecting these anomalies is the Median Absolute Deviation (MAD), a robust statistical technique that measures variability while resisting the influence of outliers. By applying MAD to the sensors' data, the framework successfully identifies significant deviations in speed that are marked as potential anomalies. This detected anomaly could point to an unexpected event, such as a vehicle malfunction or an external obstacle. The flexibility of the MAD method allows the threshold for anomaly detection to be adjusted, enabling both major and minor anomalies to be identified. This ensures that users can tailor the system to specific requirements, improving the detection of critical safety concerns or operational irregularities.

Additionally, camera data has been integrated into the analysis to further validate and understand such anomalies. Object detection and tracking methods, using YOLOv7 [23], have been employed to identify and track vehicles, pedestrians, and shuttles within the environment. These camera-based observations, when synchronized with the sensors' data, help explain the causes of detected anomalies, such as the presence of other vehicles or objects influencing the shuttle's behaviour. For example, a sudden deviation at a certain timestamp has been detected in the autonomous vehicle sensors' dataset, highlighting a significant drop in speed. This detected anomaly coming from the sensors' data has been combined with the corresponding video available from the external camera and has pointed to the reason behind the unexpected event, i.e. a vehicle overtaking the autonomous shuttle.

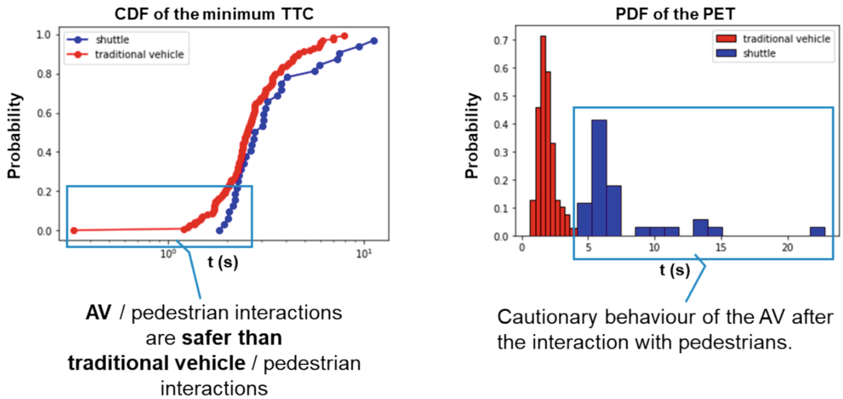

Fig. 5. Cumulative Distribution Function of the minimum TTC (left) and Probability Density Function of the PET (right): comparison autonomous shuttle (blue) vs. traditional vehicles (red) (Color figure online).

The study of AV behaviour in real traffic, especially in interactions with traditional vehicles and pedestrians, has unveiled nuanced patterns. This analysis was conducted using video data from an unsignalized crosswalk. The analysis focused on two surrogate safety measures (SSMs): Time-to-Collision (TTC) and Post-Encroachment Time (PET). These metrics were applied to assess conflicts between pedestrians and both types of vehicles (traditional and autonomous ones). A total of 33 AV-pedestrian interactions and 135 traditional vehicle-pedestrian interactions were analysed (main results displayed in Fig. 5). For TTC, which measures the time remaining before a potential collision, no statistically significant difference was found between the two types of vehicles. However, the tails of the distributions showed clear differences. In particular, pedestrian-human vehicle interactions exhibited a lower tail, indicating more dangerous situations compared to AVs. The PET, which measures the time elapsed between the vehicle and pedestrian passing through the same point, showed significant differences between the two cases. AV-pedestrian interactions had higher PET values, indicating a safer interaction compared to traditional vehicles. This is attributed to AVs stopping earlier and resuming movement more cautiously, increasing the overall safety margin during these interactions. In conclusion, while TTC did not reveal substantial overall differences, PET demonstrated that AVs create safer conditions for pedestrians by adopting more conservative driving behaviours. These findings suggest that AVs are more cautious in conflict scenarios, but they also highlight the potential for AVs to impact road network capacity due to their cautious nature. The results underline the importance of considering vehicle type in safety assessments and the potential of AVs to improve pedestrian safety in urban environments.

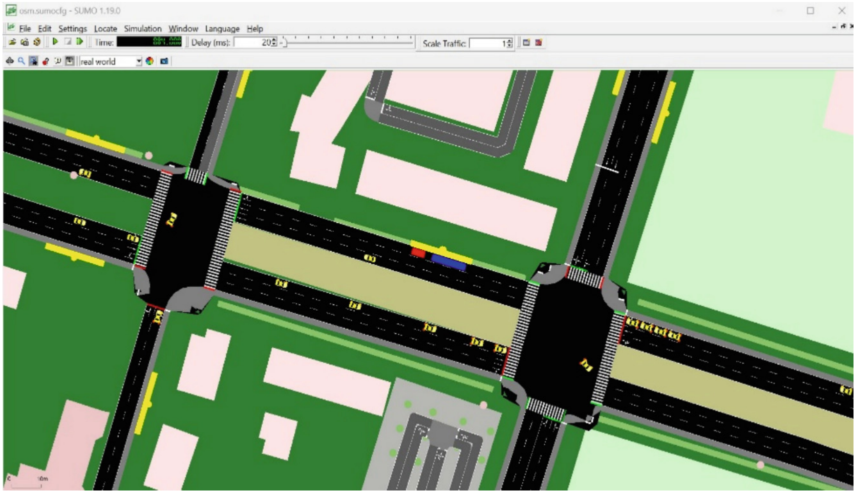

Fig. 6. Screenshot of SUMO environment in the case study.

As key outcomes of the traffic simulation phase of the proposed framework, the first result is the development of a 'baseline' to test further traffic management strategies: the testing area of the autonomous shuttle has been simulated in SUMO (Fig. 6), reproducing its real behaviour according to the sensors' data collected and analysed in the steps before. Shuttles' movement data derived from recorded shuttle data (step 1 of the framework) added into the simulation by TraCI, along pre-planned flows of passenger vehicles, public transportation, and simulated autonomous vehicles. The integration of recorded real-world data provides a high degree of authenticity and is valuable for assessing real-world scenarios and the impact of interventions or changes in traffic management with respect to the baseline, to be used as a reference scenario. Different traffic scenarios and control strategies have been tested and implemented, but -for reasons of space- they are not discussed in detail in this article and will be presented in subsequent publications. More recorded data from other real vehicles in addition to the shuttle dataset in future steps can increase the realism of the simulations.

Finally, as a result of the integrated framework, a simulation environment has been created to enable the virtual testing of autonomous vehicles. This environment is fed by sensor data collected in the field during real CCAM experiments, enriched with additional data from the surrounding environment, including infrastructure and interactions with other road users.

The simulation outcomes demonstrate the versatility of the microsimulation environment in assessing the impacts related to autonomous services, safety, and overall traffic dynamics. Optimal strategies in various contexts have been identified, contributing to the development of informed traffic management strategies. The utilization of HPC machines has significantly reduced computational time, making the simulation process efficient and scalable.

The integration of SUMO traffic microsimulation with CARLA autonomous driving simulator marks a significant achievement. Complex and realistic traffic scenarios have been provided to CARLA based on observed data, ensuring additional and customized scenarios tailored to the simulator's requirements. The co-simulation environment, operating on HPC machines, enables real-time analysis and simulation runs, reducing computational time and enhancing the overall efficiency of the testing process.

The cost-effective development process and accelerated deployment of AV demonstrations underscore the practical advantages of this integrated approach. The scalability and transferability of insights gained from the simulation environment ensure that the outcomes are applicable across diverse urban contexts.

6 Conclusions

The proposed framework is a comprehensive and versatile tool for advancing CCAM technologies in urban mobility. By leveraging real-world data, advanced machine learning techniques, and integrated simulation environments, the study contributes significantly to the understanding of AV impacts on urban traffic. The framework not only supports virtual testing of AV trials, but also provides actionable insights for urban planners and policymakers. As the development of autonomous technologies progresses, the findings and methodologies from this research will play a critical role in guiding the responsible integration of AVs into urban landscapes, ensuring both safety and efficiency in the transition to automated mobility systems.

The availability of real-world data is a vital aspect that enhances the robustness of the proposed framework. Unlike many previous studies that rely heavily on theoretical models or estimated data, this research integrates high-frequency AV data collected during real experimentations, which includes parameters such as speed profiles, interactions with pedestrians, and AV sensors' readings. The richness of this dataset allows for a more precise and realistic simulation of traffic scenarios involving AVs, making the results more applicable to real-world situations. This continuous data flow also ensures that the model remains adaptive and can be updated as more data is collected from ongoing trials.

The study's emphasis on data visualization through the PowerBI dashboard stands out as another key feature. This tool offers stakeholders a comprehensive, user-friendly interface to monitor vehicle dynamics, sensor data, and fleet-wide operations. Notably, this dashboard facilitates both real-time and offline data analysis, enabling users to gain a deeper understanding of vehicle behaviour in different traffic conditions. Such visualizations enhance decision-making processes by presenting data in a digestible format, which is particularly beneficial for stakeholders who may not have technical expertise in data analysis.

The application of machine learning algorithms for anomaly detection represents a major advancement in understanding AV behaviour. By focusing on identifying abrupt speed changes and other outlier behaviours, the framework ensures that potentially dangerous events are flagged early. This not only enhances safety but also provides valuable insights for improving AV operational efficiency. An important aspect of the anomaly detection system is its flexibility, allowing users to adjust thresholds based on specific operational needs, thus customizing the system for different urban settings.

From a traffic management perspective, the integration of the SUMO (Simulation of Urban Mobility) platform with real-world AV data creates a powerful microsimulation environment. The study examines how AVs interact with traditional traffic, including both human-driven vehicles and pedestrians. Through this simulation, the framework provides answers to crucial questions about how AVs will affect traffic flow, safety, and road network capacity in the future. By simulating various traffic scenarios may be explored, which helps urban planners and policymakers understand how the introduction of autonomous technologies might impact daily traffic patterns.

The adaptability and scalability of the proposed framework to different urban contexts are particularly noteworthy. While this framework was developed and tested using data specific to Turin, Italy, its modular design allows for adaptation to other urban environments with varying infrastructure, regulations, and traffic dynamics. Each phase of the framework, from data gathering to traffic microsimulation, is designed to be flexible and can be customized based on available data and the unique characteristics of the target city. For example, the Data Gathering phase can integrate local traffic data, AV sensor information, and infrastructure details from any city, allowing urban planners to substitute Turin's data with locally sourced inputs. The Traffic Microsimulation phase can then be configured to reflect local traffic laws, intersection designs, and typical driver behaviours by adjusting parameters to align with the new context. Moreover, the framework's reliance on widely used open-source tools, such as SUMO, facilitates adjustments for different traffic scenarios and control strategies, which can be adapted without extensive reconfiguration. Researchers wishing to replicate this framework can follow the structured methodology presented here and adapt each phase to their own data sources and computational resources, ensuring flexibility without sacrificing consistency. Future extensions of this work will aim to provide even more standardized datasets and modular code to facilitate reproducibility across diverse research environments. In addition to the technical aspects addressed in this framework, broader behavioural, regulatory, and ethical factors play a crucial role in the successful integration of AVs into urban mobility systems. From a behavioural perspective, the interactions between AVs and human road users—such as pedestrians, cyclists, and human-driven vehicles—pose unique challenges. These interactions are complex and can impact both the safety and public acceptance of AV technology. Future expansions of this framework could incorporate behavioural models to simulate and analyse how human road users react to AVs in various scenarios, which would improve the understanding of these dynamics and support more effective deployment strategies.

Regulatory challenges also affect the feasibility and scalability of AV deployment. Compliance with local traffic laws, adherence to safety standards, and adaptation to specific regulatory environments are all essential for real-world implementation. Different regions may have varying requirements for AV testing and deployment, which necessitates flexibility in the framework to accommodate such differences. Addressing regulatory requirements early in the deployment planning can help to streamline the approval process and ensure safer and more compliant AV integration. Finally, ethical considerations, such as data privacy and equitable access to AV technology, are paramount. The collection and use of AV data, especially in urban environments, must be handled with care to protect individual privacy and prevent unauthorized data usage. Additionally, ensuring that AV technologies are accessible and beneficial to all societal groups —rather than favouring specific demographics—will be vital for creating a truly inclusive mobility solution. These considerations highlight the importance of developing AV systems responsibly, with a focus on safety, inclusivity, and societal benefit.

Several challenges and limitations emerged during the study. A primary challenge was the computational demand required to run high-resolution, real-time traffic simulations, particularly when simulating complex scenarios involving interactions among diverse road users. To address this, high-performance computing (HPC) resources were employed, but the dependency on such resources may limit the applicability of the framework in contexts lacking similar computational capabilities. Future improvements could explore ways to optimize the framework to function efficiently on standard computing setups. Another limitation concerns the specificity of the data sources used in this study, which were tailored to the unique urban landscape and traffic conditions of Turin. Although the framework is designed to be adaptable to different urban contexts, its accuracy and relevance depend significantly on the availability and quality of local data, such as detailed traffic signal timings and high-frequency AV sensor data. In cities with less detailed traffic infrastructure or limited access to AV trial data, modifications to the data gathering and analysis phases would be required, potentially affecting the scalability and transferability of the framework's findings. Finally, ensuring accurate representation of human-driven vehicle behaviours and interactions with AVs remains challenging in a simulated environment. Capturing the nuances of human driving behaviour, including unpredictability and varied responses to AVs, would require additional refinements and possibly the integration of more advanced behavioural models. These challenges underscore the need for continuous development to enhance the framework's versatility and robustness across diverse urban mobility settings.

Acknowledgments. The authors would like to thank the collaborators of the Politecnico di Torino for their precious support: Prof. Silvia Chiusano, Prof. Marco Bassani, Prof. Francesco Deflorio, and Andrea Avignone. This paper is partially supported and funded by European Union's Horizon 2020 Research and Innovation Programme under Grant Agreement No 875530, project SHOW (SHared automation Operating models for Worldwide adoption), and by Horizon Europe Innovation Actions Programme under Grant Agreement No 101076791, project IN2CCAM (Enhancing Integration and Interoperability of CCAM eco-system).

References

1. SHOW project Homepage. https://show-project.eu/. Accessed 12 Sept 2024
2. IN2CCAM project Homepage. https://in2ccam.eu/. Accessed 12 Sept 2024
3. ToMove project Homepage. https://torinocitylab.it/to-move/. Accessed 12 Sept 2024
4. Anderson, J.M., Kalra, N., Stanley, K.D., Sorensen, P., Samaras, C., Oluwatola, T.A.: Autonomous Vehicle Technology: A Guide for Policymakers. RAND Corporation, Santa Monica (2016). https://www.rand.org/pubs/research_reports/RR443-2.html

5. Guériau, M., Dusparic, I.: Quantifying the impact of connected and autonomous vehicles on traffic efficiency and safety in mixed traffic. In: 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), pp. 1–8. IEEE (2020)
6. Rios-Torres, J., Malikopoulos, A.A.: Impact of connected and automated vehicles on traffic flow. In: 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC), pp. 1–6. IEEE (2017)
7. Yu, H., Tak, S., Park, M., Yeo, H.: Impact of autonomous-vehicle-only lanes in mixed traffic conditions. Transp. Res. Rec. **2673**(9), 430–439 (2019). https://doi.org/10.1177/0361198119847475
8. Nemoto, E.H., Issaoui, R., Korbee, D., Jaroudi, I., Fournier, G.: How to measure the impacts of shared automated electric vehicles on urban mobility. Transp. Res. Part D: Transp. Environ. **93**, 102766 (2021). https://doi.org/10.1016/j.trd.2021.102766
9. Anund, A., et al.: Lessons learned from setting up a demonstration site with autonomous shuttle operation – based on experience from three cities in Europe. J. Urban Mob. **2**, 100021 (2022). https://doi.org/10.1016/j.urbmob.2022.100021
10. Antonialli, F.: International benchmark on experimentations with Autonomous Shuttles for Collective Transport. In: 27th International Colloquium of Gerpisa, Paris, France (2019). https://centralesupelec.hal.science/hal-02489797v2
11. Thorhauge, M., Fjendbo, J.A., Rich, J.: Effects of autonomous first-and last mile transport in the transport chain. Transp. Res. Interdisc. Perspect. **15**, 100623 (2022). https://doi.org/10.1016/j.trip.2022.100623
12. Alvarez, L.P., Meister, W.: Microscopic traffic simulation using SUMO. In: Proceedings of the 21st International Conference on Intelligent Transportation Systems (ITSC) (2018). https://doi.org/10.1109/ITSC.2018.8569938
13. Raju, N., Farah, H.: Evolution of traffic microsimulation and its use for modeling connected and automated vehicles. J. Adv. Transp. 2444363, 29 (2021). https://doi.org/10.1155/2021/2444363
14. Lu, Q., Tettamanti, T., Hörcher, D., Varga, I.: The impact of autonomous vehicles on urban traffic network capacity: an experimental analysis by microscopic traffic simulation. Transp. Lett. **12**(8), 540–549 (2020). https://doi.org/10.1080/19427867.2019.1662561
15. SUMO Homepage. https://sumo.dlr.de/docs/index.html. Accessed 12 Sept 2024
16. Parkin, J., Crawford, F., Flower, J., Alford, C., Morgan, P., Parkhurst, G.: Cyclist and pedestrian trust in automated vehicles: an on-road and simulator trial. Int. J. Sustain. Transp. **17**(7), 762–774 (2023). https://doi.org/10.1080/15568318.2022.2093147
17. Papadoulis, A., Quddus, M., Imprialou, M.: Evaluating the safety impact of connected and autonomous vehicles on motorways. Accid. Anal. Prev. **124**, 12–22 (2019). https://doi.org/10.1016/j.aap.2018.12.019
18. Dosovitskiy, A., Ros, G., Codevilla, F., López, A., Koltun, V.: CARLA: an open urban driving simulator. In: Proceedings of the 1st Conference on Robot Learning (CoRL 2017), Mountain View, United States (2017). https://doi.org/10.48550/arXiv.1711.03938
19. Li, P., Kusari, A., LeBlanc, D.J.: A novel traffic simulation framework for testing autonomous vehicles using SUMO and CARLA. CoRR, abs/2110.07111 (2021). https://doi.org/10.48550/arXiv.2110.07111
20. Yilmaz-Niewerth, S., Häbel, R., Friedrich, B.: Developing a comprehensive large-scale co-simulation for replication of automated driving in urban traffic scenarios. Transp. Res. Procedia **78**, 522–529 (2024). ISSN: 2352-1465. https://doi.org/10.1016/j.trpro.2024.02.065
21. Shi, Y., Liu, Z., Wang, Z., Ye, J., Tong, W., Liu, Z.: An integrated traffic and vehicle co-simulation testing framework for connected and autonomous vehicles. IEEE Intell. Transp. Syst. Mag. **14**(6), 26–40 (2022). https://doi.org/10.1109/MITS.2022.3188566
22. Navya Shuttles. https://navya.tech/en/solutions/moving-people/self-driving-shuttle-for-passenger-transportation/. Accessed 12 Sept 2024

23. Wang, C.Y., Bochkovskiy, A., Liao, H.Y.M.: YOLOv7: trainable bag-of-freebies sets new state-of-the-art for real-time object detectors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7464–7475 (2023)
24. Bewley, A., Ge, Z., Ott, L., Ramos, F., Upcroft, B.: Simple online and realtime tracking. In: 2016 IEEE International Conference on Image Processing (ICIP). IEEE (2016). https://doi.org/10.1109/ICIP.2016.7533003