Traffic simulation with human in the loop: roundabout scenario in a driving simulator

Giorgio Previati[1], Gianpiero Mastinu[1], Elena Campi[1] and Massimiliano Gobbi[1]

(1) Politecnico di Milano, via La Masa 1, 20156 Milan, Italy, e-mail: {giorgio.previati, gianpiero.mastinu, elena.campi, massimiliano.gobbi}@ (polimi.it)

Abstract - *Traffic simulators are powerful tools for the simulation of vehicle interactions in many traffic conditions, but they employ approximated driver models, so the reliability of the acquired information is limited. To overcome such limitations, this paper presents a co-simulation between a widely diffused open source traffic simulator and a high end driving simulator. The coupling of the two simulators is used to investigate the effect of an actual human being driving in a simulated mixed traffic situation, where both traditional and connected and automated vehicles (CAVs) are involved. The selected reference scenario is a three-legged single lane roundabout. The behaviour of the autonomous vehicles in the simulated environment is controlled through a previously trained reinforcement learning policy. The objective is to minimize the time needed to go through the proposed roundabout. Preliminary tests are realized considering a panel of drivers on the driving simulator and with different percentages of autonomous vehicles in the simulation (20% and 80%). Results seem to indicate that the behaviour of CAVs can be easily accepted by human drivers. These outcomes deserve a further extensive statistical investigation for a final assessment.*

Keywords: *Traffic simulation, human in the loop, CAVs, roundabout, reinforcement learning.*

Introduction

Microscopic traffic simulators are powerful tools for the simulation of vehicle interactions in many traffic conditions. This allows to study traffic dynamics in most situations, including the presence in the scenario of both human-driven and automated vehicles, as done for example by (Zhong, et al., 2020). Connected and automated vehicles (CAVs) can have a positive impact on traffic flow, reducing delays (Beza, Maghrour Zefreh, and Torok, 2022) and bringing benefits to energy efficiency thus lowering pollutant emissions (Tate, et al., 2018). In order for CAVs to be safely deployed in a real environment, their interaction with traditional vehicles has to be understood. However, traffic simulators often are based on approximated driver models which cannot fully catch the complexity of vehicle behaviour (Hasan, et al., 2021), thus limiting the reliability of the information that can be acquired.

To overcome such limitations, in this paper a co-simulation between SUMO (Simulator of Urban Mobility), a widely diffuse open source traffic simulator (Lopez, et al., 2018) and a high-end driving simulator is proposed. In this way, it is possible to let a real human driver to navigate into the simulated scenario, and to assess its effect on the considered traffic scenario. This approach has been developed as a use case of AI@EDGE (AI@Edge-Consortium, 2020), a comprehensive research funded by European Commission, focused on how 5G networks can be improved by means of artificial intelligence (AI) and edge computing. This particular use case, within the broader project objectives, has the aim of understanding the requirements, benefits and limitations of edge computing and artificial intelligence applied at the task of traffic management.

The reference traffic scenario is a roundabout, negotiated by automated vehicles, simulated human-driven vehicles and a vehicle controlled by an actual human driver in the driving simulator. Roundabouts are currently one of the most critical scenarios for automated driving (Garcia Cuenca, et al., 2019a), as they represent bottlenecks in the road network. In fact, in order for autonomous driving to be considered safe and reliable, collision risk must be kept at the minimum level possible (Garcia Cuenca, et al., 2019b; Rodrigues, et al., 2018). At present, automated cars often cannot navigate this kind of infrastructure.

In literature, two main control approaches for managing automated vehicles have been attempted (Martin-Gasulla and Elefteriadou, 2021). The first one is identified as vehicle-to-infrastructure (V2I) communication, in which a central controller collects all the information from the vehicles included in the communication range, makes choices based on these data and gives instructions back to the vehicles. Conversely, applying the second approach, the decision-making process is entrusted directly to the vehicles involved in the intersection area, and it is based on the data which are exchanged among them. This configuration is referred to as vehicle-to-vehicle (V2V) communication. Both of the aforementioned approaches have shortcomings when referred to the considered scenario. V2I drawbacks are related to the infrastructure costs and to the complexity of the data management, while V2V is subjected to inefficiencies due to multi hop communication between distant vehicles. To overcome such limitation, a new approach has been developed exploiting the findings of the AI@Edge project. The new approach is based on vehicle-to-network-to-vehicle (V2N2V) communication, which relies to MEC (multi-access edge computing)/Edge nodes and exploits the features of the 5G network in terms of quality of service and low latencies enabling this application. More details on the architecture of the connect compute platform used in the AI@EDGE project are available in (AI@Edge-Consortium, 2022).

CAVs navigating the roundabout are controlled by means of a multi-agent reinforcement learning (RL) algorithm. In a RL framework, one or more agents are trained to reach a specific goal interacting with the environment (Bertsekas and Tsitsiklis, 1996). For each step or the training process, the agent performs and action on the environment, modifying it. The environment gives back to the agent a reward, associated to the fitness of the taken action with respect to the set goal. Higher rewards are given to more effective actions, allowing the agent to learn which are the best actions to reach the final objective at any state of the environment. Considering a traffic scenario, some examples of state can be the length of the queue (Abdulhai, Pringle, and Karakoulas, 2003; El-Tantawy, Abdulhai, and Abdelgawad, 2013), the positions of the vehicles and their speed (Gao, et al., 2017; Liang, et al., 2018). As reward functions, the length of the queue (Li, Lv, and Wang, 2016), the average delay (Arel, et al., 2010; Pol and Oliehoek, 2016), the cumulative delay (Liang, et al., 2018; Mousavi, Schukat, and Howley, 2017), or a combination of them (Wei, et al., 2018) are possible choices. In a multi-agent configuration, some or all the agents in the environment are interacting among each other. So it possible to coordinate their behaviour by means of a cooperative objective, thus improving the results.

The paper is structured as follows: in the next section, a description of the driving simulator is given; after that, the traffic scenario is presented; then the scheme of the co-simulation is shown; finally, the results of the preliminary tests are highlighted and some conclusion are drawn.

Dynamic driving simulator

In order to assess the safety and the acceptability of the RL behaviour policy controlling the automated vehicles, the use of a driving simulator with a human-in-the-loop is a viable option for different reasons. In fact, it allows to investigate the interaction between traditional vehicles ad CAVs in a safe environment, eliminating the risks of traffic congestions, or even worse of collisions between vehicles. Moreover, the costs related to the materials involved in the test are significantly reduced: making use of a virtual traffic environment, it is possible to modify the number of vehicles involved and their behaviour directly acting on the parameters which control the simulation. This allows to take into account different configurations without additional costs. Lastly, using a simulated scenario guarantees the repeatability of the tests, since the influence of disturbing factors such as variations in the weather conditions or in traffic flow is eliminated.

The preliminary tests are carried on at DriSMi laboratory of Politecnico di Milano (Milano, 2022), using a DiM400 Dynamic Driving Simulator. It is a cable-driven driving simulator, shown in Fig. 1, which guarantees a fully immersive experience for the driver.

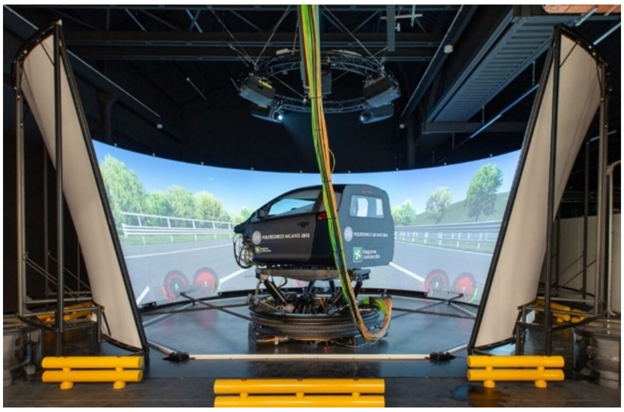

Figure 1: Driving simulator at DriSMi laboratory of Politecnico di Milano

From a graphical point of view, it is equipped with a 270°-wide 120 Hz screen surrounding the cockpit. The simulation environment is created using VI-WorldSim (VI-Grade, 2022b), which provides a 3D visualization of the traffic scenario thanks to Unreal Engine. The environment is also visible in the rear-view mirrors, realized by displays.

Figure 2: Detail of cockpit interior

The cockpit accommodating the driver is depicted in Fig. 2. Active seat belts and brake are used to reproduce the sensations related to a heavy braking manoeuvre, while the air cushions inside the seat recreate the feeling of a turning manoeuvre on the body. Five speakers are used to replicate the sounds and noises typically heard while driving, coming from both inside and outside the vehicle. The motion of the cockpit is obtained by a multi-stage system with redundant degrees of freedom. The cable-driven platform performs in-plane motions (longitudinal, lateral and yaw) at a relative low frequency (up to 3 Hz), while supporting a Steward platform which allows all six degrees of freedom, realizing smaller motions at higher frequency (up to 30 Hz). NVH (noise and vibration harshness) frequencies are reproduced by eight shakers up to 200 Hz and placed in engine and suspension mounting points.

The most important features of the driving simulator of Politecnico di Milano are listed in Tab. 1. Further details about the driving simulator can be found in (Previati, Mastinu, Gobbi, et al., 2022).

Table 1: Driving simulator data

Physical quantity	Values
Platform size	6 m x 6 m
Visual system (H)	270°
Visual system (V)	90°
Degrees of freedom	9
Longitudinal acceleration of the base	1.5 g
Lateral acceleration of the base	1.5 g
Vertical acceleration of the cockpit	2.5 g
Longitudinal travel	4.2 m
Lateral travel	4.2 m
Vertical travel	$\pm$ 298 m
Yaw angle	$\pm$ 62°
Roll angle	$\pm$ 15°
Pitch angle	$\pm$ 15°

The movements of the human-driven vehicle inside the simulation environment are computed by using a 14-degrees of freedom model present in VI-CarRealTime (VI-Grade, 2022a).All processes are controlled and synchronized by a real-time server, on which a real-time database is updated every 1 ms, corresponding to the simulation step.

Traffic scenario

The considered scenario is a three-leg single-lane roundabout developed in SUMO. In order for the co-simulation between the traffic simulator and the driving simulator to run correctly, it is fundamental to ensure the alignment of the road networks of the two simulations. This alignment can be obtained by exporting the roundabout network created in SUMO to Mathworks Roadrunner (Mathworks, 2022) and then converted to the VI-WolrdSim format via Unreal Engine (Epic-Games, 2022). The resulting scenario, shown in Fig. 3, provides the driver with a realistic view of the scene he/she is navigating.

Figure 3: Roundabout scenario

To study the interaction between CAVs and human-driven vehicles, a mixed traffic condition is simulated. In order to have a fairly traffic situation, forty vehicles are considered in the scenario. All of the vehicles present in the simulation, except for the one human-driven vehicle driven by the human inside the driving simulator, are controlled by SUMO. Traditional vehicles are reproduced using the IDM (intelligent driver model) algorithm, while CAVs behaviour is controlled by means of a RL policy. The RL policy has been trained on the same roundabout scenario. During the training phase, the traffic was constituted by forty vehicles, whose behavior was modelled by means of IDM car-following model. The parameters used in this phase have been selected based on the work of (Jayawardana and Wu, 2022), and are calibrated to reproduce the human-like driving behaviour. The policy has been trained offline and then implemented to be tested on the driving simulator. The RL algorithm has the objective to optimize the traffic flow in terms of minimizing the time needed by an autonomous vehicle to go through the roundabout. The results obtained vary according to the market penetration rate (MPR) of the CAVs, i.e. the percentage of the total traffic constituted by automated vehicles. As an indicator of the effect of MPR on the situation, the average time needed by a vehicle to go from any entrance to any exit has been measured. In a simulation with n vehicles v_1, v_2, ..., v_n, an entering timestamp $t_{v_i}^{in}$ and an exiting timestamp $t_{v_i}^{out}$ are associated to each vehicle. The average time is calculated as

$$\mu = \frac{\sum_{i=1}^{n}(t_{v_i}^{out} - t_{v_i}^{in})}{n} \quad (1)$$

The outcomes are shown in Tab. 2.

Table 2: Driving simulator data

AV percentage	$\mu(t)$ [s]
10	6.33
20	6.27
30	6.18
40	6.11
50	5.88
60	5.26
70	4.73
80	4.72
90	4.69
100	4.66

Co-simulation

The scheme for the integration between SUMO and the driver simulator is shown in Fig.4.

In order for the co-simulation between the driving simulator and SUMO scenario to run smoothly, all of the processes are integrated with the real-time server. The real-time database is used to synchronize all involved processes. The human-in-the-loop on the driving simulator drives the car by means of the pedals and of the steering wheel inside the cockpit. These signals are fed to the real-time database, and successively to VI-CarRealTime which runs the fourteen degrees of freedom model of the vehicle with a fixed simulation step of 1 ms. As output, the model gives the states of the human-driven vehicle, which are used to control the dynamic of the driving simulator. The state of the human-driven vehicles are also exploited by VI-WorldSim to generate the visual and the audio feedback the driver perceives inside the cockpit.

The SUMO simulation runs on a dedicated workstation. It is connected to the real-time database through

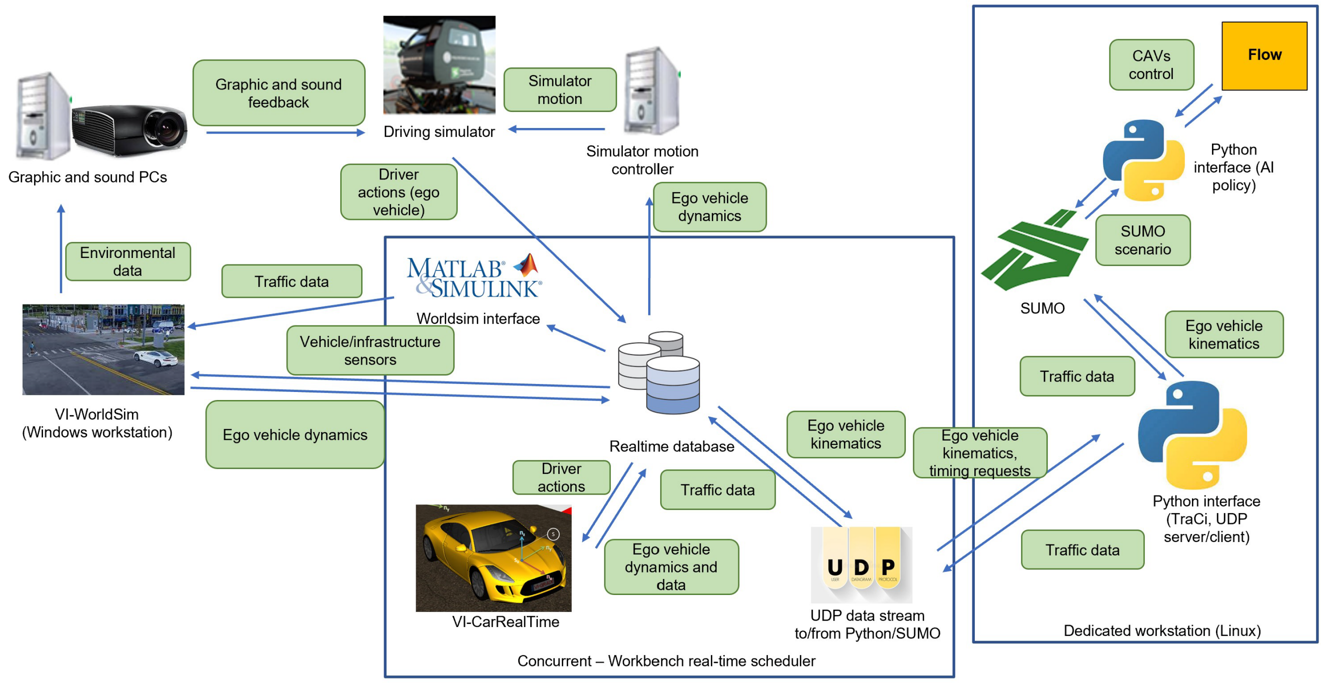

Figure 4: Scheme of the connections to the real time database

a UDP connection. The real-time sever uses this connection to send the states of the human-driven vehicle every 5 ms, thus allowing its position to be updated in the SUMO environment. The communication between SUMO and the UDP connection is realize by a Python instance exploiting the TraCI library. The UDP communication is also used as trigger to run each simulation step of 5 ms in SUMO. After each simulation step, the position of all the other vehicles, computed by SUMO, is sent to the real-time database via the UDP connection and read by VI-WorldSim, through a Matlab/Simulink interface, and used to update the virtual environment. By this scheme, the two simulations are synchronized by the real-time database through the UDP connection.

A particularly critical aspect of the co-simulation is the communication delay. The main objective of the co-simulation is to enable a real human being to drive within a simulated traffic environment. To provide the driver with a realistic experience, both simulations must run in real-time and remain synchronous. To achieve these objectives, the two simulations are synchronized with each other and with real-time using a real-time scheduler set at a fixed integration time of 5 ms. This integration time represents the allotted time for computations and proves sufficient for SUMO to receive the state of the human-driven vehicle, compute the traffic state and transmit this information to the virtual environment of the driving simulator. Due to this co-simulation scheme, there exists a fixed delay of 5 ms between SUMO and the driving simulator environment, which is imperceptible to the human driver.

Finally, a second Python interface is used to connect SUMO to the reinforcement learning policy in charge of driving the CAVs in the simulation. This connection is realized by the Flow library Flow-project, 2019.

Trajectories

SUMO is a microscopic traffic simulator, thus the trajectories of the vehicles do not consider vehicle dynamics and appear not realistic when directly used in a virtual environment for human interaction. To increase the immersivity of the driving simulation environment, the trajectories of the vehicles computed by SUMO are slightly modified before being send to the real time database. An auxiliary network has been created containing the modified trajectories of the vehicles and a correspondence table is generated relating every position im the original network to a corresponding position in the auxiliary network. During the simulation, for each position of the vehicles computed by SUMO in its network, the corresponding position in the auxiliary network is found and sent to VI-WorldSim via the real-time database. From a computational point of view, this process is very easy and quick, but it allows a more natural graphical representation of the trajectories, so that the driver perceives them as smooth and realistic. Referring to the sampling time, as SUMO samples the motion of the vehicles in the simulation at a frequency of 200 Hz, and that the screens of the driving simulator are at 120 Hz, there is no need for interpolation between the positions computed in two consecutive instants of the SUMO simulation.

Preliminary tests

The considered reference scenario, comprising the three-legged single-lane roundabout, can be navigated in a relatively short time. In order for the participants to have a more time to interact with the traffic of the other vehicles in the roundabout, they have been asked to repeat the test entering in the roundabout three times, one for each leg. Once inside the roundabout, they have been asked to leave it at the second exit, thus crossing a flow of vehicles entering the roundabout. In this way, drivers can observe the behaviour of the vehicles approaching the roundabout. For each entry leg, the manoeuvre has been repeated two times, to compare two different percentages of CAVs, i.e. 20% and 80%. Human drivers inside the simulator were not aware of the percentages of CAVs they were dealing with. They were asked to indicate the possible differences in the perception of the two traffic situations, in terms of felt safety and of traffic smoothness. The results are reported in the Tab. 3, Tab. 4 and Tab. 5.

Table 3: Perceptions related to traffic smoothness

Preferred scenario	Number of answers
20% CAVs	4
80% CAVs	6
No difference	2

Table 4: Answers associated to safety feeling

Preferred scenario	Number of answers
20% CAVs	3
80% CAVs	3
No difference	6

Table 5: Data on the global preference of the drivers

Preferred scenario	Number of answers
20% CAVs	3
80% CAVs	6
No difference	6

Regarding traffic smoothness, the data collected indicate a preference towards the scenario in which 80% of vehicles where autonomous. This means that the majority of the drivers were able to perceive a traffic flow improvement when the MPR increases. As far as it concerns the perception of safety while driving in the scenario, most of the participants did not feel any difference, so they seem to easily accept the presence of CAVs on the road. Overall, the drivers seem to accept autonomous cars in the traffic situations they faced, giving promising insights on the introduction of connected vehicles in the traditional traffic. Such conclusions should be investigated further as the panel of drivers in this preliminary test was quite limited.

Conclusions

In the paper a co-simulation between a microscopic traffic simulator and a high end driving simulator has been shown. The co-simulation allow to include a human in the loop in a traffic simulator as to understand the reactions of the human in a mixed traffic condition, in which both traditional vehicles and autonomous ones are navigating into a roundabout. Tests are performed comparing two different percentages of CAVs in the traffic, namely 20% and 80%. Their behaviour is controlled by means of a reinforcement learning policy, trained with the objective to minimize the time needed by an autonomous vehicle to cross the infrastructure.

Thanks to the driving simulator and the human-in-the-loop, it has been possible to study the perception of the driver in the considered situation. Drivers where not informed about the percentage of CAVs they were dealing with. The results obtained during the preliminary tests show that participants did not seem to be bothered by the presence of autonomous vehicle. On the contrary, there is a slight trend to prefer the situation with 80% MPR, in particular in terms of perceived traffic smoothness. This gives encouraging insights for the future of introduction of connected vehicles among traditional human-driven ones. The panel of drivers involved in the preliminary tests was quite limited, so the outcomes will be further investigated considering a wider sample.

Acknowledgements

This work is related to the research activities performed within the AI@EDGE project. This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101015922.

Bruno Lepri, Gabriele Santin and Massimiliano Luca from FBK are acknowledged for the support.

References

Abdulhai, B., Pringle, R., and Karakoulas, G. J., 2003. Reinforcement learning for true adaptive traffic signal control. *Journal of Transportation Engineering*, 129(3), pp. 278–285.

AI@Edge-Consortium, 2020. *The AI@EDGE H2020 Project*. Available at: ¡https://aiatedge.eu/¿ [Accessed Jan. 10, 2023].

— 2022. *D4.1: Design and initial prototype of the AI@EDGE connect-compute platform*. Available at: ¡https://aiatedge.eu/wp-content/uploads/2022/03/AI@EDGE_D4.1_Design-and-initial-prototype-of-the-AI@EDGE-connect-compute-platform_v1.pdf¿ [Accessed June 28, 2023].

Arel, I., Liu, C., Urbanik, T., and Kohls, A. G., 2010. Reinforcement learning-based multi-agent system for network traffic signal control. *IET Intelligent Transport Systems*, 4(2), pp. 128–135.

Bertsekas, D. and Tsitsiklis, J. N., 1996. *Neuro-dynamic programming*. Athena Scientific.

Beza, A. D., Maghrour Zefreh, M., and Torok, A., 2022. Impacts of different types of automated vehicles on traffic flow characteristics and emissions: a microscopic traffic simulation of different freeway segments. *Energies*, 15(18), p. 6669.

El-Tantawy, S., Abdulhai, B., and Abdelgawad, H., 2013. Multiagent reinforcement learning for integrated network of adaptive traffic signal controllers (MARLIN-ATSC): methodology and large-scale application on downtown Toronto. *IEEE transactions on Intelligent transportation systems*, 14(3), pp. 1140–1150.

Epic-Games, 2022. *Unreal Engine*. Available at: ¡https://www.unrealengine.com/en-US¿ [Accessed Dec. 23, 2022].

Flow-project, 2019. *Flow*. Available at: ¡https://flow.readthedocs.io/en/latest/index.html¿ [Accessed Apr. 12, 2023].

Gao, J., Shen, Y., Liu, J., Ito, M., and Shiratori, N., 2017. Adaptive traffic signal control: Deep reinforcement learning algorithm with experience replay and target network. *arXiv preprint arXiv:1705.02755*.

Garcia Cuenca, L., Puertas, E., Fernandez Andres, J., and Aliane, N., 2019a. Autonomous driving in roundabout maneuvers using reinforcement learning with Q-learning. *Electronics*, 8(12), p. 1536.

Garcia Cuenca, L., Sanchez-Soriano, J., Puertas, E., Fernandez Andres, J., and Aliane, N., 2019b. Machine learning techniques for undertaking roundabouts in autonomous driving. *Sensors*, 19(10), p. 2386.

Hasan, M., Perez, D., Shen, Y., and Yang, H., 2021. Distributed microscopic traffic simulation with human-in-the-loop enabled by virtual reality technologies. *Advances in Engineering Software*, 154, p. 102985.

Jayawardana, V. and Wu, C., 2022. Learning eco-driving strategies at signalized intersections. In: *2022 European Control Conference (ECC)*. IEEE, pp. 383–390.

Li, L., Lv, Y., and Wang, F.-Y., 2016. Traffic signal timing via deep reinforcement learning. *IEEE/CAA Journal of Automatica Sinica*, 3(3), pp. 247–254.

Liang, X., Du, X., Wang, G., and Han, Z., 2018. Deep reinforcement learning for traffic light control in vehicular networks. *arXiv preprint arXiv:1803.11115*.

Lopez, P. A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flotterod, Y.-P., Hilbrich, R., Lucken, L., Rummel, J., Wagner, P., and Wiessner, E., 2018. Microscopic Traffic Simulation using SUMO. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2575–2582. https://doi.org/10.1109/ITSC.2018.8569938.

Martin-Gasulla, M. and Elefteriadou, L., 2021. Traffic management with autonomous and connected vehicles at single-lane roundabouts. *Transportation research part C emerging technologies*, 125, p. 102964.

Mathworks, 2022. *RoadRunner*. Available at: ¡https://it.mathworks.com/products/roadrunner.html¿ [Accessed Dec. 23, 2022].

Milano, P. di, 2022. *DriSMi - Driving Simulator Politecnico di Milano*. Available at: ¡https://www.drismi.polimi.it/¿ [Accessed Dec. 23, 2022].

Mousavi, S. S., Schukat, M., and Howley, E., 2017. Traffic light control using deep policy-gradient and value-function-based reinforcement learning. *IET Intelligent Transport Systems*, 11(7), pp. 417–423.

Pol, E. Van der and Oliehoek, F. A., 2016. Coordinated deep reinforcement learners for traffic light control. *Proceedings of learning, inference and control of multi-agent systems (at NIPS 2016)*, 8, pp. 21–38.

Previati, G, Mastinu, G, Gobbi, M, et al., 2022. Influence of the Inertia Parameters on a Dynamic Driving Simulator Performances. In: *81st Annual International Conference on Mass Properties*. Society of Allied Weight Engineers, pp. 1–14.

Rodrigues, M., McGordon, A., Gest, G., and Marco, J., 2018. Autonomous navigation in interaction-based environments - A case of non-signalized roundabouts. *IEEE Transactions on Intelligent Vehicles*, 3(4), pp. 425–438.

Tate, L., Hochgreb, S., Hall, J., and Bassett, M., 2018. Energy efficiency of autonomous car powertrain.

VI-Grade, 2022a. *VI-CarRealTime*. Available at: ¡https://www.vi-grade.com/en/products/vi-carrealtime/¿ [Accessed Dec. 23, 2022].

— 2022b. *VI-WorldSim*. Available at: ¡https://www.vi-grade.com/en/products/vi-worldsim/¿ [Accessed Dec. 23, 2022].

Wei, H., Zheng, G., Yao, H., and Li, Z., 2018. Intellilight: A reinforcement learning approach for intelligent traffic light control. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2496–2505.

Zhong, Z., Lee, E. E., Nejad, M., and Lee, J., 2020. Influence of CAV clustering strategies on mixed traffic flow characteristics: An analysis of vehicle trajectory data. *Transportation Research Part C: Emerging Technologies*, 115, p. 102611.