

E-CoDrive: A Co-Simulation Framework for Testing Energy-Critical Driving Scenarios

Manfredi Napolitano
Università degli Studi di Napoli
Federico II
Naples, Italy

Alessandra Somma
Università degli Studi di Napoli
Federico II
Naples, Italy

Alessio Gambi
Austrian Institute of Technology
Vienna, Austria
The Italian Institute of AI for Industry
Turin, Italy

Andrea Stocco
Technical University of Munich
fortiss GmbH
Munich, Germany

Nicola Mazzocca
Università degli Studi di Napoli
Federico II
Naples, Italy

Abstract

Autonomous driving research has largely focused on safety while giving limited attention to non-functional aspects such as energy consumption and sustainability. As Autonomous Electric Vehicles (AEVs) become increasingly common in urban traffic, understanding how complex traffic dynamics influence their energy consumption is paramount to test whether AEVs can complete trips before battery depletion. To support energy-aware scenario-based testing of AEVs, we present E-CoDrive, a framework for reproducible closed-loop driving co-simulations that integrates an energy-consumption model, a micro-traffic simulator, and a high-fidelity driving simulator to test AEV software stacks in urban scenarios. This tool paper describes the architecture of E-CoDrive and demonstrates its applicability by testing an Autoware-based AEV stack. Our evaluation shows that varying traffic conditions produce substantial differences in vehicle energy consumption. The artifact is publicly available at <https://doi.org/10.6084/m9.figshare.32244783>, and a screencast showing the tool is available at <https://youtu.be/yX9fWHqCvgc>.

1 Introduction

Research on autonomous driving has made substantial progress in the validation of safety-critical behaviors, perception pipelines, and motion-planning logic [6, 13, 15]. However, investigating how complex environments and dynamic traffic affect energy consumption remains underexplored, despite its central role in the design of Autonomous Electric Vehicles (AEVs). Consequently, validating AEVs using existing tools that focus on safety but are oblivious to energy consumption might fail to expose misbehaviors such as idling, repeated acceleration, or detours that lead to premature battery exhaustion. Studying the interplay among traffic, energy consumption, and recovery requires coordinating traffic generation, autonomous decision-making, and energy-aware measurements in a consistent and reproducible setting.

Unfortunately, existing technologies offer limited support for integrating heterogeneous simulators and for energy-aware scenario-based testing. Driving simulators such as CARLA provide high-fidelity vehicle and environment modeling [4], while SUMO [11], CommonRoad [1], and OpenSBT [14] support scalable traffic generation and scenario construction. However, these tools are often

used in isolation [8, 17]. Although they can be integrated [9] and extended with detailed energy models [10], existing solutions primarily focus on data exchange and offer limited support for coordinated execution [7]. As a result, they may not ensure synchronized execution across traffic, vehicle dynamics, and autonomy modules. This limits their ability to accurately study how variations in traffic induce emergent effects in energy consumption [13, 15].

To address this gap, we designed E-CoDrive, a novel framework for energy-aware scenario-based testing of AEVs. E-CoDrive provides a unified orchestration layer that maintains a shared simulation timeline, coordinates state exchange across heterogeneous components, and manages the end-to-end simulation workflow. This is instrumental to enable the testing of AEVs by generating and executing energy-aware driving scenarios through synchronized closed-loop co-simulation across traffic, autonomy, and physics, and exposing energy-critical behaviors. Our preliminary evaluation shows that, compared with free-flow driving, traffic scenarios executed in E-CoDrive consistently trigger distinct and more variable energy-consumption patterns due to stop-and-go behavior.

By complementing existing research on safety-oriented testing with the analysis of non-functional properties related to energy consumption, this paper makes the following main contributions: **E-CoDrive**. A co-simulation framework that integrates traffic generation (SUMO), high-fidelity vehicle and environment simulation (CARLA), and an autonomous driving stack (Autoware Mini) within a unified orchestration layer.

Scenario Generation. The framework supports the systematic generation of energy-critical driving scenarios, including configurable traffic conditions, vehicle behavior, and energy models.

Reproducible Evaluation. The implementation leverages open-source technologies and enables reproducible execution of scenarios to evaluate how traffic conditions affect energy consumption.

2 E-CoDrive Architecture

E-CoDrive builds complex simulation scenarios in which AEVs operate under varying traffic conditions. It integrates traffic generation, high-fidelity vehicle and environment simulation, and an autonomous driving stack within a unified orchestration layer. Thanks to E-CoDrive, developers create, execute, and analyze scenarios to study how traffic dynamics influence the AEV's energy demand.

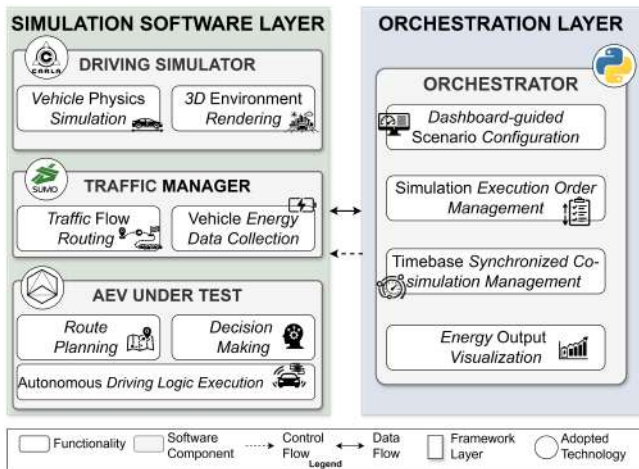


Figure 1: High-level architecture of E-CoDrive.

Figure 1 illustrates the high-level architecture of E-CoDrive. The framework is organized into two main layers. The first is the *simulation software layer*, which encompasses the core simulation tools, including the traffic management system (SUMO), the high-fidelity driving simulator (CARLA), and the AEV under test (Autoware Mini). The second is the *orchestration layer*, which comprises the orchestrator components that coordinate and synchronize interactions among these tools and the components for setting up and visualizing the energy-aware scenarios.

Note that, although the current implementation relies on specific technologies, the framework is modular and configurable, enabling the integration of alternative simulators and autonomous driving stacks. The chosen components reflect their broad adoption and popularity within the autonomous driving community.

Driving Simulator. CARLA provides 3D environment rendering and vehicle-physics simulation in which both the ego AEV and the surrounding traffic are embedded. Therefore, CARLA serves as the high-fidelity environment in which SUMO-generated traffic and Autoware Mini-issued commands are executed [4].

Traffic Manager. SUMO implements the traffic management and route generation. We opted to use it because SUMO offers flexible traffic-demand generation for heterogeneous traffic flows and enables the production and collection of vehicle-level energy data when paired with suitable electric-vehicle models [10, 11].

Ego Vehicle. The AEV under test is controlled by Autoware Mini [16], which provides the autonomous driving software stack for motion planning and vehicle control. We adopt Autoware Mini because it offers a lightweight, modular, and easily integrable autonomy stack [9]. In particular, it represents a Python-based reimplementation of the Autoware stack [5], while still providing realistic deployment characteristics, as it has been deployed on a Lexus RX450h platform at the University of Tartu [2].

Orchestrator. It is the central integration component in E-CoDrive. The orchestrator bridges interactions between users and the underlying simulation layer, hiding from them the complexity of coordinating and configuring executions. Through a dashboard, the

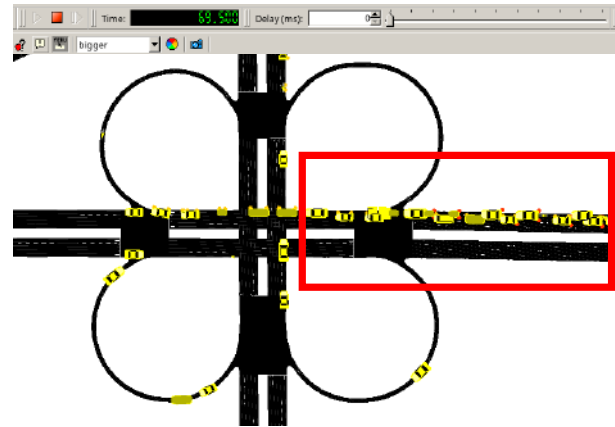


Figure 2: E-CoDrive: Congested-road scenario in SUMO.

user builds scenarios by defining the traffic conditions to be reproduced, including the amount and composition of SUMO-controlled vehicles, the relevant roads in the scenario, and the time window in which vehicles are spawned. Users can also specify the ego’s AEV energy model (e.g., MMPEVEM [10]) and driving task, including the vehicle’s static parameters, the initial battery level, and the intended route’s starting and destination points. After collecting the user configuration, invoking the required routing and scenario-generation steps, and launching the different tools, the orchestrator sets up the synchronized co-simulation. It does so by configuring the interfaces through which SUMO, CARLA, and Autoware Mini advance the simulation according to a common fixed time base. Once the simulation ends, the orchestrator collects the simulation data and visualizes it. Particularly interesting for AEV developers are the visualization of the main energy-related metrics and the driving statistics (e.g., average, minimum, and maximum speed).

3 E-CoDrive Workflow: An Example Scenario

This section illustrates E-CoDrive’s workflow using a working example. First, the user selects an available map and launches CARLA. The current implementation relies on CARLA v. 0.9.13 to ensure compatibility with Autoware Mini and offers *Town01*, *Town04*, and *Town05* maps, as these are available for both CARLA and SUMO.

Second, the user defines the traffic surrounding the ego vehicle by specifying start and destination positions and vehicle class¹ for all SUMO-controlled vehicles, thereby creating the intended traffic patterns (e.g., a traffic congestion in a selected part of the map, as highlighted by the red frame in Figure 2). This step relies on SUMO route-generation tools, in particular *duarouter*, which generates `custom_traffic.rou.xml`, a file containing all traffic routes. The orchestrator then builds `custom.sumocfg`, the SUMO configuration file used to load the map, routes, and energy output. The generated scenario is then loaded into SUMO and prepared for co-simulation.

Third, the user configures the energy model of the AEV under test by selecting one of the models, i.e., Energy [12] or MMPEVEM [10], that are currently available in SUMO.² These models estimate the

¹The vehicle class can be fixed or randomly sampled from the CARLA Vehicle Catalogue available at https://carla.readthedocs.io/en/latest/catalogue_vehicles/

²<https://sumo.dlr.de/docs/Models/Emissions.html>

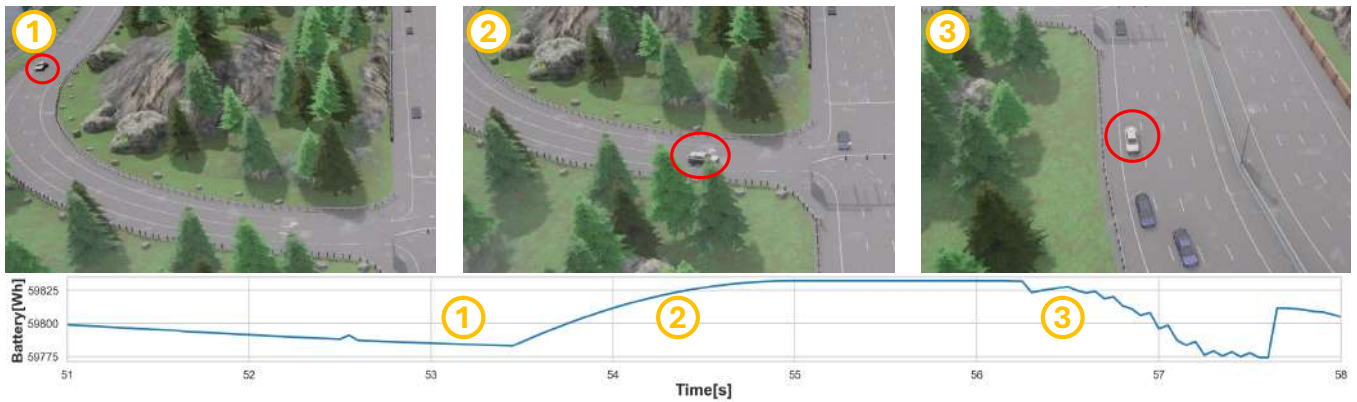


Figure 3: Representative simulated scenario frames highlighting three energy-relevant states of the AEV.

energy consumption of electric vehicles as they traverse the map during the simulation. Based on this choice, the orchestrator generates the static vehicle parameters required by the selected energy model, including battery capacity, vehicle mass, propulsion and recuperation efficiencies, and aerodynamic and rolling coefficients. To quickly explore different scenarios related to energy consumption, the user also sets the battery’s initial charge level and a threshold below which the battery is considered depleted. This step generates the file `egovtype.xml`, which SUMO uses to associate the selected energy model with the AEV under test, and the file `vtypes.json`, which CARLA uses to select the corresponding vehicle models.

Fourth, the user selects the start and destination position defining the AEV driving task, and launches Autoware Mini with the corresponding vehicle configuration and route. This step uses the `vtypes.json` file generated previously to link the AEV to the correct vehicle model.

Finally, the co-simulation starts with SUMO managing the traffic, Autoware Mini driving the AEV under test, and CARLA synchronizing the execution on a fixed time base (i.e., 20 Hz). As output, this step produces `battery.out.xml`, a file that reports instantaneous and cumulative energy consumption, speed, acceleration, and recuperation, including energy recovered during braking events.

4 Preliminary Evaluation

This section presents a preliminary evaluation of our tool, which aims to show that different traffic conditions affect the driving pattern and energy behavior of the AEV under test. In particular, we consider scenarios that induce congestion near the AEV’s destination, thus exposing variations in consumed energy, regenerated energy, and speed. Figure 3 illustrates representative traffic-induced effects on energy consumption at three timestamps, which correspond to free-flow, moderate congestion, and heavy congestion conditions. We highlight the AEV in the figure using a red circle. In frame ①, the AEV proceeds without obstacles ahead, and the battery-state profile evolves approximately linearly. In frame ②, the vehicle approaches the intersection and decelerates, leading to energy recovery through regenerative braking. Finally, in frame ③, the AEV queues in traffic, and the battery-state profile becomes irregular, potentially indicating an energy-related misbehavior. This

example illustrates the relationship between traffic conditions and the energy quantities monitored during simulation.

Building on the pipeline described above, we implemented three application scenarios defined on the CARLA maps adopted by the current implementation. For each map, the AEV route is defined by selecting the two farthest road edges, thereby creating a long mission that we execute under varying traffic conditions. In each map, we define three traffic levels, denoted as *low*, *medium*, and *high*, corresponding to 10, 25, and 50 vehicles surrounding the AEV under test, respectively. In all cases, the generated traffic is configured to congest the last segment of the AEV’s route. We configured the AEV with an energy model that approximates a Tesla Model 3 AWD profile, using a nominal battery capacity of 75 kWh. Each run starts from a fixed initial charge level of 80% (i.e., 60 kWh), rather than from a fully charged battery, to preserve plausible headroom for regenerative braking as described by Chandak et al. [3].

We compare the resulting executions against a baseline scenario in which the AEV is the only vehicle present in the network (i.e., free-flow driving). The analysis, therefore, focuses on the variation induced by traffic, examining differences in total energy consumed and regenerated, actual battery capacity, and average vehicle speed relative to the baseline. The energy metrics jointly capture traffic-induced energy exchanges and the resulting net battery depletion, while the average speed indicates whether these variations coincide with degraded driving conditions. Table 1 reports these metrics for each scenario traffic (*low*, *medium*, and *high*), map (*Town01*, *Town04*, and *Town05*), and the baseline run.

Across the three maps, the baseline scenario consistently exhibits the lowest total energy consumption. In general, the gap relative to the baseline increases as traffic density increases, suggesting that estimating energy feasibility of an AEV mission becomes more challenging under congested conditions. A small deviation from this trend is observed for *Town04* and *Town05*, where the measured energy consumption for *high* traffic is slightly lower than that for *medium* traffic. We speculate that larger maps distribute traffic more broadly, thus reducing congestion along the AEV route despite the additional vehicles present on the map.

While no strictly monotonic relationship emerges between consumed and regenerated energy across all configurations, the results

Table 1: Energy and speed metrics for each traffic scenario and setting.

	Total energy consumed/regen (Wh) ↓				Actual battery capacity (Wh) ↑				Average speed (m/s) ↑			
	baseline	low	medium	high	baseline	low	medium	high	baseline	low	medium	high
Town01	928/646	1594/1207	2457/2067	2598/2132	59721	59614	59609	59587	9.77	5.73	7.30	3.69
Town04	522/233	5585/4911	5221/4574	4394/3806	59711	59327	59353	59413	12.81	11.48	10.61	8.59
Town05	254/107	920/720	2056/1766	1391/1156	59853	59801	59709	59765	13.19	12.74	12.63	6.45

consistently show that traffic conditions produce energy behaviors that are substantially different from the baseline scenario. Compared with free-flow driving, traffic scenarios generally exhibit higher energy consumption and regeneration, along with lower average speed, indicating more frequent acceleration and braking phases due to congestion and queuing.

Overall, these results suggest that traffic has a non-negligible impact on the AEV's energy demand. In the most pronounced cases, the increase in battery depletion reaches approximately 137 – 358 Wh over routes spanning only 1.3 – 2.1 km. Although preliminary, these findings indicate that the interaction between traffic conditions and battery consumption is worth further systematic investigation. In particular, denser traffic appears to induce more frequent acceleration, braking, and speed adaptations, which may substantially affect short-term energy usage and overall route efficiency. These observations motivate future studies on larger and more diverse traffic scenarios to better understand the relationship between traffic dynamics and energy consumption in AEVs.

5 Conclusions and Future Work

This paper presented E-CoDrive, a co-simulation framework for studying how traffic conditions affect the energy behavior of AEVs. By integrating SUMO, CARLA, and Autoware Mini into a synchronized, guided execution workflow, we support reproducible scenarios, coordinated simulations, and the collection of energy measurements under closed-loop conditions. The preliminary evaluation shows that traffic dynamics affect the AEV energy profile, battery depletion, and speed relative to the free-flow baseline. This supports traffic-aware energy validation as a relevant non-functional testing objective for AEVs.

Our future plans include extending E-CoDrive to include automated, search-based testing capabilities for battery-critical scenarios. Instead of manually defining traffic conditions, search-based algorithms can systematically explore the scenario space to identify combinations of traffic and energy conditions that lead to problematic behaviors. This is particularly relevant because the number of parameters defining such scenarios is very large. Our goal is to automatically generate scenarios in which an AEV that would normally complete its route under free-flow traffic instead experiences delays, stop-and-go behavior, or rerouting conditions that push the battery close to depletion. Such scenarios could help identify energy- and traffic-aware failures, including unsafe immobilization, inefficient routing decisions, or behaviors that negatively affect surrounding traffic. In the long term, this could support the development and evaluation of battery-aware planning and fallback strategies.

6 Data Availability Statement

A replication package with our framework, evaluation scripts, and datasets is available at <https://doi.org/10.6084/m9.figshare.32244783>.

References

- [1] Matthias Althoff, Markus Koschi, and Stefanie Manzing. 2017. CommonRoad: Composable benchmarks for motion planning on roads. In *2017 IEEE Intelligent Vehicles Symposium (IV)*. 719–726. doi:10.1109/IVS.2017.7995802
- [2] Autonomous Driving Lab, University of Tartu. 2023. Vehicle. <https://adl.cs.ut.ee/lab/vehicle>. Accessed: 2026-03-06.
- [3] Gaurav A. Chandak and A. A. Bhole. 2017. A review on regenerative braking in electric vehicle. In *2017 Innovations in Power and Advanced Computing Technologies (i-PACT)*, 1–5. doi:10.1109/IPACT.2017.8245098
- [4] Alexey Dosovitskiy, Germán Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. 2017. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning (CoRL)*.
- [5] Autoware Foundation. 2022–2025. Autoware Core/Universe. <https://github.com/autowarefoundation/autoware>.
- [6] Yuan Gao, Mattia Piccinini, and Yuchen et al. Zhang. 2026. Foundation Models in Autonomous Driving: A Survey on Scenario Generation and Scenario Analysis. *IEEE Open Journal of Intelligent Transportation Systems* (2026).
- [7] Cláudio Gomes, Casper Thule, David Broman, Peter Gorm Larsen, and Hans Vangheluwe. 2018. Co-simulation: A survey. *Comput. Surveys* (2018).
- [8] Kai Huang, Ziheng Zhang, Xiang Wang, Yiheng Tao, and Zhiyuan Liu. 2025. Life-cycle carbon emissions of autonomous electric vehicles in varying traffic situations. *Transportation Research Part D: Transport and Environment* 146 (2025). <https://www.sciencedirect.com/science/article/pii/S1361920925002810>
- [9] Shinpei Kato, Shota Tokunaga, Yuya Maruyama, Seiya Maeda, Manato Hirabayashi, Yuki Kitsukawa, Abraham Monrroy, Tomohito Ando, Yusuke Fujii, and Takuya Azumi. 2018. Autoware on Board: Enabling Autonomous Vehicles with Embedded Systems. In *2018 ACM/IEEE 9th International Conference on Cyber-Physical Systems (ICCP)*. 287–296. doi:10.1109/ICCP.2018.00035
- [10] Lucas Koch, Dominik S. Buse, Marius Wegener, Sven Schoenberg, Kevin Badalian, Falko Dressler, and Jakob Andert. 2021. Accurate physics-based modeling of electric vehicle energy consumption in the SUMO traffic microsimulator. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*.
- [11] Daniel Krajzewicz, Jakob Erdmann, Michael Behrisch, and Laura Bieker. 2012. Recent Development and Applications of SUMO – Simulation of Urban MObility. In *International Journal on Advances in Systems and Measurements*.
- [12] Tamás Kurczveil, Pablo Álvarez López, and Eckehard Schnieder. 2014. Implementation of an Energy Model and a Charging Infrastructure in SUMO. In *Simulation of Urban Mobility*. Springer Berlin Heidelberg, 33–43.
- [13] Guannan Lou, Yao Deng, Xi Zheng, Mengshi Zhang, and Tianyi Zhang. 2022. Testing of autonomous driving systems: where are we and where should we go?. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE 2022)*.
- [14] Lev Sorokin, Tiziano Munaro, Damir Safin, Brian Hsuan-Cheng Liao, and Adam Molin. 2024. OpenSBT: A Modular Framework for Search-based Testing of Automated Driving Systems. In *Proceedings of ICSE-Companion '24*. ACM.
- [15] Shuncheng Tang, Zhenya Zhang, Yi Zhang, Jixiang Zhou, Yan Guo, Shuang Liu, Shengjian Guo, Yan-Fu Li, Lei Ma, Yinxing Xue, and Yang Liu. 2023. A Survey on Automated Driving System Testing: Landscapes and Trends. *ACM Trans. Softw. Eng. Methodol.* 32, 5, Article 124 (July 2023), 62 pages. doi:10.1145/3579642
- [16] Autonomous Driving Lab University of Tartu. 2026. Autoware Mini. https://github.com/UT-ADL/autoware_mini/tree/release/nodes/detection/lidar/cluster
- [17] Ying Zhang, Jinchao Chen, Tao You, Yingjie Zhang, Zhaohua Liu, and Chenglie Du. 2024. Energy-Aware Optimization of Connected and Automated Electric Vehicles Considering Vehicle-Traffic Nexus. *IEEE Transactions on Industrial Electronics* 71, 1 (2024), 282–293. doi:10.1109/TIE.2023.3245204