

EV-sim: Urban Traffic and Battery Dynamics Simulator for Degradation-Prognostics and Range-Aware Decision-Making for Electric-Vehicle Operations

Jorge E. García Bustos¹, Benjamín Brito Schiele², Bruno Masserano³, Cristóbal E. Allendes⁴, Ricardo Salas-Espiñeira⁵, Fernando Oropeza Suárez⁶, Catalina Platz-Gamboa⁷, Lukas Gleisner J.⁸, Diego Troncoso-Kurtovic⁹, Francisco Jaramillo-Montoya¹⁰, Heraldo Rozas¹¹, and Marcos E. Orchard¹²

^{1,3,4,5,6,7,8,9,10,11,12} *Department of Electrical Engineering,
Faculty of Physical and Mathematical Sciences, University of Chile, Santiago, Chile*

*jorgegarcia@ug.uchile.cl
bruno.masserano@ug.uchile.cl
cristobal.allendes@ug.uchile.cl
ricardo.salas.e@ug.uchile.cl
fernando.oropeza@ug.uchile.cl
catalina.platz@ug.uchile.cl
lukas.gleisner@ug.uchile.cl
diego.troncoso.k@ug.chile.cl
francisco.jaramillo@uchile.cl
heraldo.rozas@ug.uchile.cl
morchard@ing.uchile.cl*

² *Intelligent System Prognostics Group, Aerospace Structures and Materials Department,
Faculty of Aerospace Engineering, Delft University of Technology, Delft, 2629HS, The Netherlands
bbritoschiele@tudelft.nl*

ABSTRACT

Existing open-source traffic tools accurately reproduce driver behavior and congestion for conventional internal-combustion vehicles. However, in the case of electric vehicles (EVs), they often fail to incorporate critical electrical variables, such as battery voltage, power demand, and State-of-Health, which limits their applicability in operational planning and decision-making. This paper introduces a simulation platform tailored for EVs that bridges the gap between traditional transportation models and the needs of the PHM community in electromobility. The proposed platform combines power and energy consumption profiles derived from Gaussian Mixture Models with physics-based representations of battery behavior. Model parameters are calibrated using a publicly available dataset collected in Ann Arbor, Michigan. Each trip is partitioned into segments based on abrupt changes in speed, ensuring uniform operating conditions within segments and enhancing model transferability across routes. The platform simulates vehicle speed, electrical

power demand, State-of-Charge (SoC), terminal voltage, and incremental capacity loss at each simulation step. Battery degradation is estimated through an empirical model fitted to long-term cycling data. A case study demonstrates the simulator's ability to compare route alternatives between a shared origin and destination. Results show that the shortest path is not always the most energy-efficient nor the least degrading, highlighting the value of health-aware routing. The platform will be publicly released to enable reproducible testing of SoC estimation, range prediction, and degradation forecasting without requiring extensive instrumentation or prolonged field testing.

1. INTRODUCTION

In recent years, there has been a significant shift toward electromobility, particularly with the adoption of Battery Electric Vehicles (BEVs) in urban environments. This trend is expected to continue accelerating in the coming years, as governments and industries worldwide increase their investments in electric mobility technologies and supporting infrastructure (International Energy Agency, 2025). However, the full implications of this transition are not yet completely under-

Jorge E. García Bustos et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

stood, as available data often covers only a limited range of operating scenarios. This is particularly relevant considering that, while lithium-ion batteries (LIBs) offer a viable and efficient energy storage solution, they also present notable challenges. One of the most critical aspects is the uncertainty associated with a vehicle's achievable range for a given route and its State-of-Charge (SoC). Although this may initially appear to be a straightforward relationship, empirical evidence reveals that there is often no consistent correlation between SoC and the actual distance an electric vehicle (EV) can travel. This discrepancy is further compounded by the influence of the route profile, where factors such as regenerative braking and energy recovery during downhill travel can significantly alter the effective range.

Moreover, unlike their internal combustion engine counterparts, the energy storage capacity of EV batteries degrades over time, introducing an additional layer of uncertainty in range estimation. On one hand, degradation affects the operational performance of the battery through effects such as capacity fade, where the charge that can be extracted decreases over time; increases in internal resistance, which lowers maximum output current and increases heat generation; and reductions in efficiency, leading to higher energy losses (Rahman & Alharbi, 2024). For range estimation, these effects accumulate, leading to inaccurate results (Hu, Xu, Lin, & Pecht, 2020). On the other hand, degradation also affects operational safety, where high degradation levels can increase the risk of catastrophic failures such as thermal runaway. As Gao, Li, Offer, and Wang (2024) demonstrates, degradation affects the optimal safety thresholds used to generate early-alarm systems for failures, potentially leading to an increase in false negatives when degradation effects are not considered.

These two key points highlight the importance of analyzing the evolution of degradation during operation and modeling its effects on the vehicle's operational characteristics, particularly its electrical and thermal variables. Several modeling approaches have been proposed to address this challenge. For instance, Pérez et al. (2017) presents a degradation model for LIBs for erratic SoC swing ranges, utilizing data provided by manufacturers, along with similarity-based models, to calculate degradation rates from different usage profiles. Allendes et al. (2024) introduces an electrothermal model for LIBs that accounts for the influence of internal temperature on performance. In this model, temperature-dependent changes in usable capacity are incorporated, and internal resistance is linked to heat generation, both of which can impact the cell's electrical behavior (Spitthoff, Shearing, & Burheim, 2021). As Hou, Yang, Wang, and Zhang (2020) demonstrates, elevated internal temperatures can, in and of themselves, accelerate degradation.

In this context, simulation emerges as a powerful tool to

model the complex interactions between battery behavior, vehicle dynamics, and operating conditions. Simulations enable the exploration and analysis of scenarios that are not directly observable from empirical data alone, leveraging existing theoretical and experimental knowledge to model system behavior. Additionally, simulators can be useful tools for testing and benchmarking EV routing strategies. Currently, EV routing strategies are often validated in customized simulation environments (Futalef, Muñoz-Carpintero, Rozas, & Orchard, 2023), which makes it difficult to compare the performance across different models.

Although no existing solution dedicated to simulating EVs in urban environments was identified, related approaches can be found in the literature. For example, DTUMOS (Yeon, Eom, Jang, & Yeo, 2023) is a large-scale urban mobility digital twin primarily oriented toward demand-supply generation, vehicle dispatch and routing, and time-of-arrival prediction, using inputs such as trip records and OpenStreetMap data (Bennett, 2010). While effective for analyzing passenger mobility and system-level performance indicators, this simulator does not incorporate native energy or electrochemical models, nor does it provide modules for lithium-ion battery behavior or degradation. Consequently, it is not feasible to use DTUMOS to simulate EV fleets.

A different approach is offered by GAMA (Amouroux, Chu, Boucher, & Drogoul, 2009), a generic, open-source platform for agent-based modeling that supports a broad spectrum of applications, ranging from ecological studies and search-and-rescue operations to urban traffic management. Despite its flexibility, GAMA does not include built-in EV-specific energy models or frameworks for capturing lithium-ion battery degradation. As a result, it cannot adequately address the challenges unique to EV simulation, particularly those involving battery behavior, range uncertainty, and long-term degradation dynamics.

This work introduces a novel simulation platform called *EV-sim* and specifically designed for BEVs. Unlike existing simulators, *EV-sim* uniquely integrates detailed predictions of electrical variables such as voltage, current, instantaneous power, and energy consumption, providing an unprecedented depth of analysis for EV operations. A key innovation is the incorporation of a real-time battery degradation model, which accurately reflects battery capacity loss under diverse operational scenarios. Furthermore, *EV-sim* models the vehicle's usage profile through detailed power and energy consumption profiles, enabling real-time simulation of various routes and generating comprehensive statistical models. This capability allows users to evaluate the risks associated with route choices, directly linking battery degradation and usage patterns. The detailed modeling of EV-specific behaviors is achieved with remarkably low computational demands,

ensuring the simulator's practical applicability in real-time decision-making contexts and fleet management scenarios.

The paper is organized as follows: Section 2 provides a theoretical background, detailing essential concepts and prior work relevant to the development of the proposed platform. Section 3 describes the proposed simulation methodology, highlighting its distinctive features. Section 4 presents the obtained results from various case studies and simulations. Section 5 offers a detailed discussion and analysis of these results. Finally, Section 6 summarizes the main conclusions derived from this research and suggests potential directions for future improvement.

2. THEORETICAL BACKGROUND

2.1. Battery model

To analyze battery dynamics, it is important to employ an equivalent circuit model that captures the system's response under different operational conditions. In this study, a Thévenin equivalent model is adopted, incorporating a controlled voltage source to represent the open-circuit voltage (V_{oc}) along with a resistance that accounts for internal losses, for example, heat dissipation (Díaz et al., 2020). The open-circuit voltage V_{oc} is modeled as a function of the SoC, which typically follows a characteristic nonlinear trend.

The V_{oc} and SoC of a battery can be defined as:

$$V_{oc} = V + I \cdot R, \quad (1)$$

where V is the measured terminal voltage, I is the measured electrical current, and R is the internal resistance. The SoC is formulated as follows:

$$SoC(t) = SoC_0 - \frac{\int_{t_0}^t i(\xi) d\xi}{C}, \quad (2)$$

where C denotes the nominal capacity of the battery in ampere-hours (Ah), and $i(\xi)$ is the current as a function of time, integrated from the initial time t_0 to the current time t .

An analogous formulation can be used to define the State-of-Energy (SoE), which accounts for both current and voltage over time. The SoE is given by:

$$SoE(t) = SoE_0 - \frac{\int_{t_0}^t u(\xi) i(\xi) d\xi}{E}, \quad (3)$$

where E represents the nominal energy of the battery in joules (J), and $u(\xi)$ and $i(\xi)$ are the instantaneous voltage and current, respectively, integrated over the same time interval.

Building on this, Pola et al. (2015) introduced a five-parameter expression for V_{oc} , as described in Eq. 4.

$$\begin{aligned} v_{oc}(SoC_t) = & v_L + (v_0 - v_L) \cdot e^{\gamma(SoC_{t-1})} \\ & + \alpha \cdot v_L(SoC_{t-1}) \\ & + (1 - \alpha)v_L(e^{-\beta} - e^{-b\sqrt{SoC_t}}) \end{aligned} \quad (4)$$




Figure 1. Battery open circuit voltage response through three distinct operational zones during the discharge cycle.

Each term in Eq. 4 captures the behavior of the open-circuit voltage (V_{oc}) curve (Figure 1) across distinct operating regions of the battery's SoC. The first exponential term, $(v_0 - v_L) \cdot e^{\gamma \cdot SoC_{t-1}}$, models the initial discharge phase, where a sharp voltage decline is observed as the battery begins to deplete from a fully charged state. This steep drop is governed by the parameter γ , which controls the curvature of the exponential response.

The second term, $\alpha \cdot v_L \cdot SoC_{t-1}$, corresponds to the mid-range discharge region. In this zone, the voltage exhibits a nearly linear decline as SoC decreases, which is a common characteristic of Li-ion batteries in practical applications. The slope of this decline is determined by the weighting parameter α , allowing the model to represent the relatively stable voltage behavior observed in this central segment.

The final nonlinear term, $(1 - \alpha) \cdot v_L \cdot (e^{-\beta} - e^{-b\sqrt{SoC_t}})$, accounts for the accelerated voltage drop observed in the later stages of discharge. This segment reflects the diminished electrochemical activity at low SoC levels, often resulting in a sharp and nonlinear fall-off in voltage. The parameter β modulates this curvature, enabling the model to accurately represent the end-of-discharge dynamics.

This five-parameter formulation offers a phenomenological yet efficient approximation of the typical Li-ion $V_{oc}(SoC)$ profile. As described by Pola et al. (2015), the simplicity of this structure supports reliable parameter identification with limited training data while still capturing the key nonlinearities across the full SoC range.

2.2. Battery Degradation Model

LIBs degrade progressively as they experience charge and discharge cycles. This cumulative effect reduces the amount of energy the battery can store. The degradation is commonly quantified through the State-of-Health (SoH), which represents the current maximum capacity of the battery relative to its original nominal capacity.

The degradation model employed in this work is based on tracking the effect of usage conditions across discrete *equivalent cycles*. Each equivalent cycle corresponds to the cumulative discharge of the battery's current full capacity, regardless of the specific sequence or irregularity of the individual charge/discharge events. This abstraction allows the model to handle non-symmetrical, erratic operational profiles more realistically than conventional full-cycle definitions.

After each equivalent cycle, the battery's SoH is updated by applying a multiplicative degradation rate, denoted as η , following the recursive relation:

$$\text{SoH}_{k+1} = \eta_k \cdot \text{SoH}_k \quad (5)$$

The degradation rate η_k is not fixed but instead depends on operating conditions experienced during the k -th cycle. Specifically, the model characterizes each equivalent cycle as a function of first, the difference between the maximum and minimum SoC values during the cycle (Swing Range, from now on defined as SR), the mean SoC value over the cycle duration (Average Swing Range, from now on defined as ASR), and the average ambient temperature during the equivalent cycle

First, the SR and ASR features are mapped to a degradation rate η_k^{SBM} using a Similarity-Based Model (SBM), as described in (Pérez et al., 2017). In this approach, degradation rates associated with a specific proportion of the total life cycle of a nominal degradation campaign (SR = 100 in this case) are linked to particular combinations of SR and ASR, and then used as reference points. Each reference point is connected to a fraction of the total equivalent cycles the battery can perform compared to the SR = 100 case, shown in Table 1 as "Life Cycle Percent".

Here, the degradation rate for a new cycle is estimated via a weighted interpolation using the 3-nearest neighbors in the feature space. This allows the model to flexibly infer η_k^{SBM} under arbitrary SoC conditions, even those not explicitly observed during training. The learned mapping effectively captures the nonlinear dependence of capacity fade on SoC dynamics, aligning with well-established electrochemical aging behavior.

Then, the relation between total life cycle and ambient temperature established in (Zhou, Qian, Allan, & Zhou, 2011) is

Table 1. Normalized degradation factors for each SR.

SR	Life Cycle Percent
100-0	1.00000
100-25	0.78750
75-0	1.12525
100-50	0.43750
75-25	0.68750
50-0	1.03125
100-75	0.40625
75-50	0.29700
62.5-37.5	0.28125
50-25	0.62500
25-0	1.00000

used to modulate the degradation rate η_k as a function of the average ambient temperature experienced by the battery during the k -th cycle. According to this study, elevated temperatures accelerate degradation due to increased reaction rates within the cell, consistent with Arrhenius-type behavior. In LIBs, high temperatures promote cell oxidation and the formation of a Solid Electrolyte Interphase (SEI), leading to irreversible capacity loss, as illustrated in Figure 2. This figure shows that the life cycle of a lithium-ion battery can be reduced by more than half when operating at 55°C compared to 25°C. To account for this, a temperature correction factor η_{temp} is defined as the ratio of the expected life cycle at nominal conditions to that at the measured average temperature. The final degradation rate for the equivalent cycle is then computed as:

$$\eta_k = \eta_k^{\text{SBM}} \cdot \eta_{\text{temp}}, \quad (6)$$

where η_k^{SBM} is the base degradation rate obtained from the previously mentioned SBM using the SoC profile, and η_{temp} adjusts this rate to reflect the additional aging effects induced by thermal exposure. This formulation ensures the model captures the nonlinear sensitivity of battery aging to environmental conditions, especially under high-temperature operation.

2.3. Maximum Driving Range

A previous study presented the Maximum Driving Range (MDR) framework. This route-segmentation methodology merges data-driven inference with an electrochemical battery model to predict the range of EVs (García Bustos et al., 2025). Rather than depending on fixed calibration factors or extensive historical logs, MDR subdivides each trip into physically uniform segments identified by roadway grade, surface quality, and typical traffic flow attributes.




Figure 2. Relation between total life cycle and operating ambient temperature for a 2000-cycle battery. [Extracted from (Zhou et al., 2011)]

Within every segment, machine-learning regressors estimate the expected speed profile together with the associated power and energy demand; these estimates then serve as inputs to a reduced-order battery model that tracks voltage dynamics, SoC evolution, and cumulative energy depletion. The interaction between the statistical and physics layers captures both driving behavior and electrochemical responses with high fidelity.

By deriving the segmentation adaptively during model execution, the framework naturally extends to routes that have never been driven before, an important capability in data-sparse contexts. By propagating uncertainty through Monte Carlo (MC) simulations, MDR produces a probabilistic map of hazard zones, defined as segments where the predicted terminal voltage nears or falls below the battery’s cut-off threshold, indicating a high likelihood of power loss and providing actionable guidance for drivers and fleet managers. MC validation reported in García Bustos et al. (2025) showed that this probabilistic characterization reliably highlights risk regions while accounting for sensor noise and stochastic variations in driving patterns.

2.4. Gaussian Mixture Model

A Gaussian Mixture Model (GMM) is a probabilistic model that approximates an unknown, potentially multi-modal density as a probability-weighted sum of K multivariate normal distributions called “components” (Reynolds, 2015). Mathematically, this model is expressed as:

$$p_{\Theta}(\mathbf{x}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x}; \mu_k, \Sigma_k) \quad (7)$$

where each component k is defined by its parameters Θ_k : the mean vector μ_k , the covariance matrix Σ_k , and the weight π_k , which represents the probability that a given data point belongs to component k . These parameters collectively define local Gaussian clusters within the dataset.

To estimate the parameters of a GMM, we begin with a set of independently and identically distributed (i.i.d.) samples $\{\mathbf{x}_n\}_{n=1}^N$. The Expectation–Maximization (EM) algorithm is then applied to maximize, in an iterative manner, the following log-likelihood function:

$$\ell(\Theta) = \sum_{n=1}^N \log p_{\Theta}(\mathbf{x}_n). \quad (8)$$

In the E-step, EM introduces binary latent variables $z_{nk} \in \{0, 1\}$. We set $z_{nk} = 1$ when sample $\mathbf{x}_n$ is generated by component k and $z_{nk} = 0$ otherwise. Because each observation is assumed to come from exactly one component, we enforce $\sum_{k=1}^K z_{nk} = 1$ for all n . These latent variables let us treat the problem as if the component assignments were known, converting it into a complete-data scenario (McLachlan & Krishnan, 2008).

The EM algorithm then alternates between two steps. In the *E-step*, it computes the posterior probability (or *responsibility*) that component k was responsible for generating data point $\mathbf{x}_n$, given the current parameter estimates Θ^{old} :

$$\gamma_{nk} = \frac{\pi_k \mathcal{N}(\mathbf{x}_n; \mu_k, \Sigma_k)}{\sum_{j=1}^K \pi_j \mathcal{N}(\mathbf{x}_n; \mu_j, \Sigma_j)}, \quad (9)$$

where γ_{nk} represents the posterior probability of the k -th latent variable and the n -th data point.

Then, in the *M-step*, the model parameters are updated based on the computed posterior probabilities. These posterior responsibility, γ_{nk} , represents the degree of belief (or expected membership) that a given data point $\mathbf{x}_n$ was generated by component k . For the case of GMMs, the mean parameter associated with each latent variable k is computed as a weighted expected value:

$$\mu_k^{\text{new}} = \frac{1}{N_k} \sum_{n=1}^N \gamma_{nk} \mathbf{x}_n, \quad (10)$$

with the covariance matrix follows:

$$\Sigma_k^{\text{new}} = \frac{1}{N_k} \sum_{n=1}^N \gamma_{nk} (\mathbf{x}_n - \mu_k^{\text{new}})(\mathbf{x}_n - \mu_k^{\text{new}})^{\top}, \quad (11)$$

where the normalizing factor N_k is calculated following:

$$N_k = \sum_{n=1}^N \gamma_{nk}, \quad (12)$$

hence ensuring that the effect of posterior responsibilities remains bounded.

This process is guaranteed to monotonically increase the log-likelihood at each iteration and converges to a local maxi-

mum. Model complexity is typically controlled by choosing K using the Bayesian Information Criterion (BIC) or a held-out log-likelihood, and by applying constraints on Σ_k (e.g., diagonal, tied, or spherical) to reduce overfitting in high-dimensional settings.

3. METHODOLOGY

The following section will present the methodology employed to develop the proposed simulator. First, the general simulation pipeline will be introduced. Then, a segmentation heuristic protocol that discretizes the urban routes will be described. Finally, this section presents a probabilistic model of energy and time consumption that defines the speeds and energy requirements of the simulated vehicle at each time step.

3.1. EV route simulation procedure

The proposed simulator consists of the following elements:

1. A graph that represents the simulated urban environment.
2. Agents that represent the EVs under study.
3. An EV energy and time consumption model.

At each step of the simulation, these elements interact with each other to sequentially update the state of the agents. Figure 3 represents how the mentioned elements interact at each simulation phase. The rest of this subsection will describe each element of the simulator.

In the initialization step of the simulation, the program generates a graph that represents the simulated urban environment. In this graph, vertices represent locations on the street and edges represent the roads that connect the locations. Note that the application of a graph transforms the spatial domain from continuous to discrete; hence, the number of vertices and edges should be large enough to represent the time constants of the battery dynamics and power requirements. The urban graph is constructed starting from an OpenStreetMap (OSM) base map. Nodes are then inserted according to the segments generated by the proposed segmentation protocol, which ensures that node and edge densities accurately reflect the street features relevant for EV simulation. Routing constraints are inherently captured through the connectivity and directionality of edges provided by the OSM data and refined during segmentation. The next subsection describes a heuristic segmentation approach that defines how vertices and edges are located to obtain a suitable graph representation of the urban environment.

Once the graph is initialized, the simulator proceeds to instantiate the fleet of agents that model the EVs under study. Each agent is fully described by a state vector with two components: current position in the map and current SoE of the battery. In addition, each agent is characterized by a set of battery parameters and a routing policy. The battery parameters

allow modeling the electrical behavior of the battery using Eqs. 1, 3 and 4. The routing policy of the agent corresponds to a function that chooses the next vertex that the agent will visit based on the available contextual information. Formally, this function can be expressed as follows in Eq. 13:

$$\pi : \mathbb{X} \rightarrow \mathbb{V} \quad (13)$$

$$\pi(x) = v_{next} \quad (14)$$

where π is the policy function, $\mathbb{X}$ is an arbitrary input space of the policy, $\mathbb{V}$ is the set of map vertices, and v_{next} is the next vertex that the agent plans to visit.

After the initialization phase, the program starts the simulation. It is important to note that, in the current simulations, vehicles do not perform any charging. While future extensions may include charging behavior, for the present study, all agents deplete their batteries without recharging. At each simulation step, the map iteratively consults which agents are available to move to the next edge of the graph. For all agents that are available to move, the following routine is executed. First, the agent consults its policy to determine which edge it will travel through to reach the next vertex. Then, the map consults an EV energy and time consumption model to define the energy requirements and the time consumed traveling through the selected edge. The definition of the EV energy and time consumption model will be explained in the following subsection.

At last, the agent is allowed to move to the next vertex along the edge. It is noteworthy to mention that once the simulator decides to move an agent, it also marks it as unable to move until the internal clock of the simulation reaches the predicted time consumption. This procedure allows coordinating the fleet agents over time while indirectly modeling the traffic conditions of each vertex. Also, once the movement is completed, the simulator updates the SoE of the agent's battery and calculates its terminal voltage using Eqs. 3 and 1 respectively. This procedure allows for monitoring the electrical behavior of the EVs during the simulation.

All agent initialization parameters—including the number of agents per run, initial SoE, battery parameters (capacity and internal resistance ranges), and routing policies—are fully configurable according to user requirements. These parameters are implemented as an external function to the simulation pipeline, allowing flexible adjustment for different experimental scenarios. Similarly, the simulation time resolution and total duration can be configured to adapt the granularity and length of each run.

Finally, the simulation stops once one of the agents runs out of charge. Authors acknowledge that this stopping condition might be overly simplistic; therefore, including vehicle charging options is proposed for future updates of the simulator. The outputs of the finished simulator are a time series




Figure 3. Proposed urban EV simulation pipeline.

of data for the SoE, voltage at terminals, and positions of all agents. This information can be useful for later evaluation of the routes taken by the agents in a post-processing stage. For instance, this work uses the battery degradation model presented in Section 2.2 to calculate the degradation rate of the battery associated with the traveled route. By computing this information, it would be possible to assess whether a routing policy is acceptable in terms of its impact on the battery lifetime.

3.2. Route Segment Protocol

This work adopts and extends a segmentation-based preprocessing framework for EV telemetry data, initially introduced in prior work (García Bustos et al., 2025). The segmentation protocol plays a crucial role in developing transferable speed and energy consumption models by partitioning raw driving data into segments characterized by relatively homogeneous environmental and operational conditions. Each resulting segment serves as an individual modeling unit during both the training and inference phases, enabling statistical generalization to previously unvisited routes.

Segmentation proceeds iteratively, terminating a segment either when the cumulative distance exceeds a predefined upper bound (100 meters by default) or when a *Focus Point* is encountered. These *Focus Points* correspond to specific markers along the route (such as traffic lights, gates, or pedestrian

crossings) that typically induce abrupt changes in vehicle velocity. Consequently, their occurrence indicates transitions in the driving context, prompting the initiation of a new segment. This dual criterion ensures that segments are both spatially coherent and behaviorally informative.

After initial segmentation, a post-processing stage is applied to refine segment definitions further. First, segments with insufficient spatial length are eliminated or merged. Specifically, adjacent segments shorter than 25 meters (a criterion not previously included) are merged to maintain statistical reliability while preserving meaningful semantic transitions. Once segments are finalized, each one is represented as a single data point by aggregating internal measurements. Numerical features (e.g., velocity, acceleration, and battery power) are summarized by computing their arithmetic mean across the segment duration. Conversely, categorical features (e.g., street type or presence of an intersection) are encoded according to specialized rules. Notably, for attributes such as “Focus Points” or “Street Type,” the most frequent non-zero value within the segment is selected, defaulting to zero only when necessary. This strategy prioritizes informative infrastructure indicators, thereby improving the overall quality of feature representation.

Formally, consider a trip divided into N segments, $\mathcal{S}_1, \dots, \mathcal{S}_N$, according to the segmentation procedure described above. Each segment n contains a set of consecutive

raw data points denoted by $\mathcal{S}_n = k_1, k_2, \dots, k_{\mathcal{T}_n}$, where $\mathcal{T}_n = |\mathcal{S}_n|$ represents the number of samples within that segment. For every numerical variable u_k present in the data set, including battery voltage, current, SoC, vehicle speed, slope, and meteorological data, the representative value at the segment level is calculated as the arithmetic mean:

$$u_n = \frac{1}{\mathcal{T}_n} \sum_{k \in \mathcal{S}_n} u_k. \quad (15)$$

Energy consumption, however, is treated differently. The dataset provides incremental energy consumption measurements from the battery at each sampling interval. Hence, the representative energy consumption for segment n corresponds to the cumulative energy drawn within that segment:

$$E_n = \sum_{k \in \mathcal{S}_n} E_k. \quad (16)$$

The proposed segmentation protocol has demonstrated its value in case studies comparing multiple routes connecting the same origin-destination pairs. By abstracting complex driving patterns into standardized, statistically consistent segments, the methodology provides a robust foundation for reproducible simulations and the deployment of predictive analytics within EV fleet operations. Furthermore, the modular design and compatibility with publicly available data facilitate its adoption in ongoing research and future developments in electromobility.

3.3. Electric-Vehicle Energy and Time Consumption Modeling

The simulation procedure described in Section 3.1 makes use of an EV energy and time consumption model to define the cost of traveling through an edge of the map in terms of energy and time. Since the urban environment is inherently stochastic, a consumption model must also represent such uncertainty in its outputs to correctly model the phenomenon (García Bustos et al., 2025). Although neural network architectures have proven to be successful for this task (García Bustos et al., 2025), these models are considered computationally intensive. Hence, the present work employed an ensemble of GMMs to compute the consumptions due to their ability to represent uncertainty and low computational cost. Note that for this application, the computational cost of models is an important factor, since they will be employed extensively during simulations.

Inspired by previous work that tackles this problem (García Bustos et al., 2025), the consumption model is trained using a two-step approach. An initial model is trained to predict EV speed from the edge features, then a second model uses these predictions to obtain energy consumption predictions.

Finally, the time consumption is calculated using spread prediction and the edge length. Mathematically, the first model can be defined as a function:

$$\mathcal{S}_{\theta_s} : \mathbb{X} \times \mathbb{R} \rightarrow [0, 1] \quad (17)$$

$$\mathcal{S}_{\theta_s}(x, s) = p(s|x), \quad (18)$$

where $\mathcal{S}_{\theta_s}$ is the speed prediction model, $x \in \mathbb{X}$ is the input feature vector, θ_s are model parameters and $p(s|x)$ is the probability of traveling with a speed s through the edge. Similarly, the second model follows:

$$\mathcal{E}_{\theta_e} : \mathbb{X} \times \mathbb{R} \times \mathbb{R} \rightarrow [0, 1] \quad (19)$$

$$\mathcal{E}_{\theta_e}(x, s, e) = p(e|x, s), \quad (20)$$

where $\mathcal{E}_{\theta_e}$ is the energy consumption prediction model, $s \in \mathbb{R}$ is a sample of the previously obtained vehicle speed prediction, θ_e are model parameters and $p(e|x, s)$ is the probability of using e units of energy to travel through the edge. Since these models characterize probabilities, values are sampled from the predicted distributions in each simulation step.

To implement the probabilistic models $\mathcal{S}_{\theta_s}$ and $\mathcal{E}_{\theta_e}$, an ensemble of GMMs was employed. In this architecture, an initial GMM is trained using the EM algorithm. Once the model is trained, it can infer a set of latent variables z_k from an input vector $x \in \mathbb{X}$ following Bayes' theorem:

$$p(z_k|x) = \frac{p(x|z_k) \cdot p(z_k)}{\sum_{k=1}^N p(x|z_k) \cdot p(z_k)}, \quad (21)$$

In the context of this work, the latent variable z_k can be interpreted as a driving profile that is conditional on the characteristics of the road.

With the computed posterior probability, a second GMM can be trained to model the output distribution from the latent variables conditioned on the input. Mathematically, this is expressed as:

$$p(y|x) = \sum_{k=1}^N p(y|z_k)p(z_k|x), \quad (22)$$

where $p(y|x)$ is the distribution of the model output y given an input x . Note that for this model, $p(y|z_k)$ is a Gaussian distribution; hence, $p(y|x)$ is also a GMM. The primary difference between the first and second GMMs is that for the first model, the latent variable z_k is still unknown; hence, it must be inferred in the E-step of the EM algorithm. In contrast, for the second GMM, the latent variable is already defined using Eq. 21, thus it is already given for the EM algorithm.

3.4. Reproducibility

This section presents the implementation details required to reproduce all results reported in this work. It includes the dataset version, data cleaning and segmentation rules, the splitting strategy used to avoid information leakage, feature processing, modeling configurations, and the environment setup. All experiments were run using fixed random seeds, and Monte Carlo simulations were performed with a consistent number of repetitions across experiments.

3.4.1. Data and splitting.

The dataset used in this study is the Vehicle Energy Dataset (VED) (Oh, Leblanc, & Peng, 2022), collected in Ann Arbor between November 2017 and November 2018. The complete dataset comprises 32,940 trips covering nearly 600,000 km from 383 privately owned vehicles instrumented at 1 Hz. The fleet includes 264 gasoline, 92 hybrid, and 27 plug-in hybrid or battery-electric units, with three 2013 Nissan Leafs contributing about 7,500 km of purely electric operation. Each EV log provides time-aligned measurements of GPS (latitude and longitude), vehicle speed, HV battery voltage and current, ambient air temperature, and auxiliary power usage. Data is organized in weekly CSV files labeled `VED_mmddyy_week.csv`, while static attributes such as engine type, drive train, and weight are stored separately.

The altitude of the road network ranges from 220 m to 330 m above sea level, with ambient temperatures between -15°C and 35°C , occasionally accompanied by rain or snow. These conditions establish a realistic framework for energy consumption and traffic modeling, serving as a rigorous benchmark for validating the proposed simulator under representative urban driving scenarios.

To prevent information leakage, we partition the data at the level of Trip and Vehicle, allocating 70 % of trips to training, 20 % to validation, and 10 % to testing. Normalization statistics (mean and standard deviation) are computed exclusively on the training set and applied consistently during validation and testing stages.

The evaluation presented in Section 4 focuses on four representative test trips from the BEV 455, with specifications detailed in Table 3 and route plots in Figure 4.

3.4.2. Models and hyperparameters.

The predictive framework uses two GMMs: a speed model that estimates segment-level speed distributions conditioned on climatic and road variables and an energy model that estimates segment-level energy consumption conditioned on speed and contextual covariates. Tables 5 and 4 summarize the predictors and targets; the energy model is conditioned on the speed model by the mean speed of the predicted segment.

Both GMMs were trained once and serialized (`joblib`) for consistent evaluation. During simulation, the trained models are loaded to generate predictions; Monte Carlo sampling is used to propagate segment-level uncertainty to route-level performance.

Table 2. Hyperparameters for the speed and energy GMMs.

Model	Components	Covariance type	Initializations
Speed model	175	full	5
Energy model	175	tied	5

3.4.3. Training, evaluation, and MC.

Both models are trained with EM using k-means seeding for component means, full covariance matrices, and an εI diagonal regularization for numerical stability. EM iterates until the absolute change in log-likelihood is below the tolerance or the maximum number of iterations is reached. When multiple initializations are specified, the solution with the highest final log-likelihood is retained. Convergence diagnostics (final log-likelihood and iteration count) are recorded.

Evaluation is conducted on the held-out test set with the same preprocessing statistics used during training. We report per-trip and overall mean absolute error (MAE) for the speed model and the energy model, and include uncertainty bands obtained via MC propagation at the route level. MC simulations draw from the fitted mixture distributions to propagate segment-level uncertainty through the energy estimates; unless otherwise stated, we use $N = 100$ runs to compute confidence intervals and risk curves.

To ensure reproducibility of stochastic components (initialization and sampling), we fix the global random seed to 42. The NumPy random generator used in seeding, sampling, and toy data generation is initialized with this seed so that k-means initialization and EM restarts are deterministic given the same configuration. All experiments are executed with the same configuration files and environment specification to guarantee repeatability across runs.

At this stage, the simulator codebase is in active development as a beta version. Our plan is to release the finalized open source implementation, along with the preprocessing and training scripts, once the framework is fully stabilized. This staged release strategy ensures that the community will access a tested and documented version that can be reliably reused.

4. RESULTS

4.1. Case Study

This case study analyzes four test trips performed using a single battery-electric vehicle. The selected trips were drawn from a publicly available driving dataset collected in Ann Ar-

bor, Michigan, between November 2017 and October 2018 (Oh et al., 2022).

Each selected trajectory was recorded using the BEV 455 (2013 Nissan Leaf), following a distinct urban route (see Figure 4), thereby exposing the simulator to varied traffic density, slopes, stop-and-go patterns, and climate conditions. Table 3 summarizes basic characteristics such as distance, cumulative elevation gain, and average ambient temperature to contextualize the data set.

Although some trips appear spatially similar, their departure times and weather profiles differ markedly. This diversity is intentional, as it allows us to assess how well the segment-level speed and energy regressors generalize when faced with traffic and climate regimes unseen during training. Importantly, the case studies reported here are designed primarily to demonstrate the simulator’s end-to-end capabilities, linking route choice, traffic, ambient conditions, and cycle-wise battery degradation, rather than to benchmark the underlying predictive models used. A broader benchmarking campaign across additional routes, weather conditions, and vehicle/pack configurations is planned as future work.




Figure 4. Geographic layout of the four evaluation routes extracted from the Ann Arbor dataset. Each colored polyline represents the recorded trajectory of a single trip.

Having established the accuracy of these statistical models, we now move to a second validation layer that examines the complete simulation pipeline. Four agents, each corresponding to a test trip, are simulated in parallel on the street network

Table 3. Key characteristics of the four test trips.

Trip	Distance [km]	Elevation change [m]	Avg. Temp. [°C]
1197	14.62	16.49	1.53
1509	13.09	-0.50	13.50
1648	15.80	-26.35	28.74
2323	14.58	24.50	9.00

graph, replicating the historical trajectories. At each segment, the agents update their position, velocity, power demand, and residual charge. Weather variables retrieved from the World Weather Online API (World Weather Online, 2025) are incorporated, exposing the simulated battery to the same ambient conditions as those experienced by the actual vehicle during the recorded trips.

The evaluation focuses on two observable quantities: net energy consumption and total travel time. The energy is obtained from the power demand model, and the travel time is calculated from the speed model. Then both outputs are compared with the values recorded on the board. (Figure 5) visualizes these comparisons and discusses simulator fidelity segment by segment.

To further illustrate the platform’s flexibility, two complementary analyses are presented. First, leveraging the integrated degradation model, the simulator estimates battery degradation by quantifying capacity loss accrued per duty cycle, as previously defined in Section 2.2. This capability enables predictive assessment of battery health over repeated execution of identical trips, highlighting implications for long-term fleet management. Second, an additional analysis evaluates the risk of battery disconnection using the MDR methodology introduced in Section 2.3 earlier. By simulating an agent under a constrained routing policy, in which the vehicle moves from a starting location to a specified destination without revisiting previously traversed nodes, the platform assesses how user-defined routing constraints influence operational risks such as prematurely reaching battery cut-off voltage.

These two analyses are based on the precision of the segment-level speed and energy estimates generated by our Gaussian-mixture regressors. The predictor sets used at each stage are summarized in Tables 4 and 5. Table 4 lists the 13 variables that drive the Stage-1 speed model, grouped into meteorology, road geometry, traffic context, and segment dynamics. The Stage-2 energy model (Table 5) relies on only five inputs, namely the predicted mean speed for the current segment, the predicted mean speed for the previous segment, the segment length, and two descriptors of weather and road conditions. By separating the feature sets in this way, the simulator first captures how external factors influence driving behavior and then converts that behavior into battery-power demand without redundant information.

Table 4. Predictor variables for the speed model (Stage 1).

Category	Variables
Meteorology	Outside-air temperature (°C); Precipitation rate (mm/h); Relative humidity (%)
Road geometry	Segment length (m); Road-grade ratio; Smoothed altitude (m); Signed slope angle (°)
Traffic context	Posted speed limit (km/h); Node type ^a ; Street class ^b ;
Dynamics	Distance since trip start (m) Segment duration (s); Speed – previous segment (km/h)

^a crossing, roundabout, *etc.*; ^b motorway, arterial, residential, *etc.*

Table 5. Predictor variables for the energy model (Stage 2).

Category	Variables
Meteorology	Outside-air temperature (°C); Precipitation rate (mm/h)
Road geometry	Segment length (m); Road-grade ratio
Dynamics	Predicted mean speed – current segment (km/h) Predicted mean speed – previous segment (km/h) Segment duration (s)

4.2. Energy and Speed Modeling

This subsection provides an integrated analysis of both the speed and energy consumption models, evaluating their accuracy across the selected trips. Table 6 summarizes the MAEs for both models. The results indicate that the speed model achieves MAEs ranging from 4.615 km/h (Trip 1197) to 6.381 km/h (Trip 1509), while the accumulated energy model reports MAEs between 0.132 kWh (Trip 1509) and 0.217 kWh (Trip 1648). In general, despite variations in route characteristics, traffic conditions, and elevation changes, the models exhibit robust performance, suggesting a high degree of generalizability.

Figure 5 illustrates the predicted accumulated energy consumption compared to the Ground Truth (GT) across each route. The predicted mean closely matches the observed data, with minor discrepancies typically associated with abrupt changes in driving conditions, such as transitions from urban to suburban environments or terrain elevation. Importantly, the probabilistic confidence intervals (CIs) effectively encapsulate the observed values, highlighting the model's efficacy in quantifying the uncertainty.

Notably, for Trip 1197 (Figure 5a)) and Trip 2323 (Figure 5c)), the accumulated energy errors remain below 0.2 kWh for the majority of segments, consistent with the low overall MAEs of 0.177 kWh and 0.166 kWh, respectively. Trip 1509 (Figure 5b)), despite having the shortest distance among these four routes, exhibits slightly larger

Table 6. MAEs for speed and accumulated energy models.

Trip ID	Speed MAE (km/h)	Accumulated Energy MAE (kWh)
1197	4.615	0.177
1509	6.381	0.132
1648	5.914	0.217
2323	5.424	0.166

uncertainty bands during the mid-trip segments (20–60), possibly due to more variable traffic patterns. Nevertheless, the predicted mean curve still aligns closely with the true cumulative energy profile. In Trip 1648 (Figure 5d)), the model captures the gradual increase in energy use when ascending and descending small hills, with the CI widening modestly toward the end of the trip.

Figure 6 illustrates the segment-level speed predictions for the same four trips. Again, the predicted mean (blue line) reproduces the key oscillations and plateaus observed in the GT (open circles). The 90 % CI (gray shading) generally encompasses observed speed values, demonstrating that the model effectively captures variability in driving behavior due to stop-and-go traffic, traffic lights, and varying speed limits.

For Trip 1197 (Figure 6a)), which traverses both urban grid and arterial highways, the model adjusts rapidly to sudden speed drops. In Trip 1509 (Figure 6b)), predictions closely follow the GT on highway segments (speed approximately 70 km/h), but show slightly larger uncertainty around low-speed segments (< 20 km/h), likely due to transient stops. Trip 2323 (Figure 6c)) features multiple stop-and-go cycles: the model's CIs widen during these cycles, but the mean curve still remains within 5 km/h of the measured speed. Finally, Trip 1648 (Figure 6d)) includes one long uphill segment where speeds drop below 30 km/h; the predicted mean underestimates slightly in that region, but maintains coverage within the CI.

Overall, the reported errors verify that the learning framework reproduces segment-level dynamics with sufficient fidelity for downstream decision-making. Speed MAEs below 6.5 km/h represent at most 9 % of the mean cruising speed observed in the dataset, while energy MAEs under 0.22 kWh correspond to relative deviations below 6 % of trip energy demand. These bounds fall comfortably within the tolerances typically required for range-aware routing, charge-planning, and preliminary degradation assessments, confirming that the stand-alone speed and energy models provide a reliable foundation for the more elaborate simulations examined in subsequent sections.

4.3. Real-World Driving Simulations for Electric Vehicles

To evaluate the simulator's ability to reproduce trip-level metrics, Tables 7 and 8 compare GT observations with predictions from *EV-sim* across four representative trips. Specif-




Figure 5. Accumulated energy consumption (kWh) for four distinct trips. Circles represent ground-truth measurements, solid lines indicate predicted means, and shaded regions denote 90 % CIs.

ically, these metrics assess total travel time, final SoC, and total energy consumption per trip.


Table 7 presents the raw simulated and observed values for each variable. For instance, Trip 1197 exhibits nearly identical SoC values (64.10 %) between *EV-sim* and the GT, while the predicted travel time deviates by less than 5 % from the actual duration. Similarly, energy estimates for Trips 1648 and 2323 differ by just 0.02 and 0.15 kWh, respectively, confirming that the simulator preserves fidelity in cumulative indicators.

Table 8 complements this analysis by reporting the MAE and associated percentage deviations. Energy MAEs remain below 0.51 kWh across all trips, while SoC MAEs stay within 1.1 percentage points. Notably, Trip 1509 shows the highest energy error at 9.16 %, likely due to transient low-speed patterns that compound power estimation inaccuracies. Conversely, Trip 1648, despite a moderate travel time error (5.13 %), yields the lowest energy deviation (0.32 %), sug-


gesting a robust match between speed, power, and voltage profiles under varying terrain.

Together, these results demonstrate that *EV-sim* provides reliable trip-level predictions of energy and time, with deviations small enough to support use cases such as energy-aware routing, charge-planning, and range forecasting under real-world conditions.


To translate aggregate trip statistics into a route-specific risk metric, we use the MDR framework to estimate the probability of battery disconnection. Because the Ann Arbor dataset does not include trips that fully deplete the pack, we synthetically concatenate repeated out-and-back traversals of a selected route. For each MC replication, the simulator draws a fresh sequence of segment-level speed and power samples, executes the round-trip loop until the terminal voltage reaches the cut-off threshold, and records the distance traveled at that instant. Accumulating these outcomes yields a cumulative-risk curve that quantifies how the likelihood of an involuntary shutdown grows with distance.




a) Trip 1197



b) Trip 1509



c) Trip 2323



d) Trip 1648

Figure 6. Vehicle speed profiles (km/h) for four distinct trips. Circles denote ground-truth speed samples, solid lines represent predicted means, and shaded regions indicate 90 % CIs.




Figure 7. The cumulative probability of vehicle disconnection obtained from 100 MC repetitions of Trip 2323.

Figure 7 shows the resulting curve for Trip 2323. The probability of disconnection remains below 10 % for the first 70 km, then rises sharply and reaches 100 % at approximately 108 km. The companion plot in Figure 8 shows the corresponding terminal-voltage envelope. The median trajectory drops from 392 V to approximately 338 V, and the 90 % band




Figure 8. Mean terminal-voltage trajectory (blue) and its 90 % confidence band (grey) for the same MC experiment.

intersects the cut-off line at the same distance, confirming internal consistency.

These results demonstrate how *EV-sim* can convert segment-level uncertainty into a practical, probabilistic range-safety metric, enabling fleet operators to proactively anticipate disconnection risk and schedule charging or rerouting actions.

Table 7. Comparison between simulated (*EV-sim*) and GT values for total travel time, final SoC, and energy consumption.

Trip ID	Total Time (s)		Final SoC (%)		Energy (kWh)	
	EV-sim	GT	EV-sim	GT	EV-sim	GT
1197	1810.72	1902.80	64.10	64.10	5.71	6.08
1509	1571.44	1521.80	55.91	56.46	3.76	3.45
1648	913.97	963.40	60.87	61.22	3.78	3.76
2323	1652.23	1674.10	84.89	85.98	4.88	5.03

Table 8. MAE and relative errors between simulation and GT data for travel time, final SoC, and energy consumption.

Trip ID	Travel Time		Final SoC		Energy Consumption	
	MAE (s)	Error (%)	MAE (%)	Error (%)	MAE (kWh)	Error (%)
1197	420.46	4.84±11.92	0.068	0.014	0.505	6.11±5.91
1509	201.25	3.26±6.59	0.554	0.981	0.485	9.16±6.97
1648	238.81	5.13±15.73	0.348	0.568	0.329	0.32±2.04
2323	289.93	1.31±13.34	1.082	1.258	0.401	2.93±8.21

4.4. Battery-health projection for vehicle 455

All the trips in the dataset that correspond to vehicle 455 were concatenated end-to-end. Every time the accumulated charge reached the full usable capacity of the pack, we declared one *equivalent cycle* and logged the SR, the ASR, and the outside-air temperature associated with that cycle. The resulting data is highly heterogeneous: SR spans from 0.3 % to 39 %, ASR from 49 % to 94 %, and ambient temperature from -8°C to 27°C . This variety makes the profile representative of the different seasons and driving missions encountered by the vehicle.

The battery currently shows a measured SoH of 90 % and the manufacturer quotes 2000 equivalent cycles to reach the 70 % end-of-life limit under 100 % swings. Starting from these conditions, we propagated the calendar forward, assuming that future use repeats the same sequence of cycles. One hundred MC runs were performed; at every cycle, the degradation model drew a random realization of the rate η conditioned on the observed SR, ASR, and temperature.

Figure 9 gives the probability density of the number of cycles required to hit the 70 % threshold, whereas the right-hand panel shows the corresponding cumulative distribution. The most likely time to the end of life is approximately 900 equivalent cycles, and there is a 90 % chance that failure will occur between 830 and 1020 cycles.

Figure 10 plots the SoH trajectory obtained by stacking the cycles in chronological order. The solid line is the ensemble mean, and the shaded band marks the 5–95 % CI.

The monotonic decline confirms that the duty cycle is sufficiently aggressive to drive the pack from 90 % to 70 % in roughly 100 ± 80 cycles. These results illustrate how the proposed degradation module converts heterogeneous field data into a quantitative, uncertainty-aware forecast of the remaining useful life.




Figure 9. Left: probability density of the number of equivalent cycles required to reach the 70 % SoH threshold. Right: corresponding cumulative distribution.

5. DISCUSSION

The combined use of GMMs for segment-level feature prediction and a physics-based battery representation yields accurate and reliable estimates of speed, energy demand, and SoC across diverse urban routes. Speed MAEs remain below 6.5 km/h and accumulated-energy MAEs stay under 0.53 kWh for trips of 10–20 km, despite pronounced variations in traffic density, slope, and stop-and-go behavior. Ninety-percent confidence bands consistently enclose most GT points, indicating that the uncertainty quantification is well-calibrated.

At the trip level, Tables 7 and 8 show that predicted travel time, final SoC, and energy consumption align closely with recorded data. SoC errors are limited to 1.1 % and energy deviations stay below 10 %, confirming that the framework captures cumulative effects such as power loss on hills and recuperation during deceleration. These tolerances meet the requirements for range-aware routing, charge-planning, and preliminary degradation assessment in fleet operations.

The extended maximum-driving-range experiment demonstrates that the simulator also produces meaningful risk metrics when empirical data do not cover full-depletion cases. By




Figure 10. Mean SoH (dark blue) with 90 % confidence band (light blue) obtained from 100 MC simulations.

looping a representative route in an MC setting, *EV-sim* constructs a cumulative probability curve of battery disconnection. For Trip 2323 the risk remains under 10 % up to 70 km and reaches 100 % near 108 km, while the companion voltage envelope confirms that failure coincides with the cut-off threshold. This ability to translate segment-level stochasticity into a route-specific safety metric enables proactive scheduling of charging stops or dynamic rerouting.

Finally, the degradation prognostic module ties these operational results to the battery's long-term outlook. For vehicle 455, the MC projection suggests that only 1000 ± 80 equivalent cycles remain before the 70 % SoH limit, even though the pack is now at 90 %. This short horizon matches the aggressive duty cycle just characterized and the wide ambient span from -7°C to 26°C , a combination known to accelerate both cycling and calendar aging.

Several limitations suggest directions for future work. First, the evaluation will be expanded to encompass a broader range of routes, climates, and traffic scenarios. Second, additional electro-thermal effects, such as temperature-dependent internal resistance and active cooling, will be incorporated to refine degradation forecasts. Third, live traffic, weather, and charging-station data will be integrated via external APIs, enabling closed-loop studies where vehicle decisions respond to real-time conditions. Addressing these points will strengthen *EV-sim* as a decision-support tool for large-scale electric-fleet management.

6. CONCLUSIONS

This work introduces *EV-sim*, a simulation framework that bridges data-driven traffic modeling with battery-physics awareness to support health-centered decision-making for electric-vehicle fleets. By emulating realistic driving behavior while tracking electrochemical states, the platform provides a laboratory-grade environment in which researchers and operators can explore how route profiles, traffic condi-

tions, and cell aging interact without the expense of continuous on-board instrumentation or long cycling campaigns.

A built-in MC module extends the analysis from point predictions to risk quantification. Segment-level uncertainty propagates to cumulative probabilities of voltage cut-off, yielding disconnection-risk curves that translate directly into proactive charge scheduling and route selection. Because the architecture is modular, additional physical layers, such as detailed thermal models or live data streams from weather and charging networks, can be integrated with minimal effort, turning *EV-sim* into a flexible bench for control strategies and prognostic algorithms.

Validation on four real urban trips shows that the framework predicts speed, energy demand, and SoC within margins suitable for everyday fleet management. Synthetic full-depletion tests further demonstrate that *EV-sim* can estimate range risk even when historical logs never reach cut-off voltage, a prerequisite for safety-critical deployment. Taken together, these results position *EV-sim* as a practical and extensible tool for research and operational planning. A forthcoming public release will enable reproducible studies across a wide spectrum of routes, climates, and vehicle types, encouraging the community to accelerate advances in health-aware electrified transport.

Planned work includes explicit modeling of charging-station interactions, integration of temperature-dependent resistance and active thermal control, and extension to mixed fleets that incorporate heavy-duty and autonomous vehicles. These upgrades will further position *EV-sim* as a reference environment for research on energy and health-aware decision support in electric mobility.

ACKNOWLEDGMENT

This work was supported in part by ANID FONDECYT 1250036, Advanced Center for Electrical and Electronic Engineering, ANID Basal Project AFB240002. The work of Jorge E. García Bustos has been supported by ANID-PFCHA/Doctorado Nacional/2022-21221213. The authors also gratefully acknowledge the **Department of Electrical Engineering, Faculty of Physical and Mathematical Sciences, University of Chile**, for financial assistance through its 2025 International Conference Grant Program.

REFERENCES

- Allendes, C., Beltrán, A., García, J. E., Troncoso-Kurtovic, D., Masserano, B., Brito Schiele, B., ... Rangarajan, S. (2024, November). Modeling and simulation of thermal effects on electrical behavior in lithium-ion cells. *Annual Conference of the PHM Society*, 16(1). doi: <https://doi.org/10.36001/phmconf.2024.v16i1.4080>
- Amouroux, E., Chu, T.-Q., Boucher, A., & Drogoul, A.

- (2009). Gama: An environment for implementing and running spatially explicit multi-agent simulations. In *Agent computing and multi-agent systems* (p. 359–371). Springer Berlin Heidelberg. doi: https://doi.org/10.1007/978-3-642-01639-4_32
- Bennett, J. (2010). *Openstreetmap*. Packt Publishing Ltd.
- Díaz, C., Quintero, V., Pérez, A., Jaramillo, F., Burgos-Mellado, C., Rozas, H., ... Cárdenas, R. (2020). Particle-filtering-based prognostics for the state of maximum power available in lithium-ion batteries at electromobility applications. *IEEE Transactions on Vehicular Technology*, 69(7), 7187–7200. doi: <https://doi.org/10.1109/TVT.2020.2993949>
- Futalef, J.-P., Muñoz-Carpintero, D., Rozas, H., & Orchard, M. E. (2023). An online decision-making strategy for routing of electric vehicle fleets. *Information Sciences*, 625, 715–737.
- Gao, X., Li, R., Offer, G. J., & Wang, H. (2024). *Mapping safety transitions as batteries degrade: A model-based analysis towards full-lifespan battery safety management*. arXiv. doi: <https://doi.org/10.48550/arXiv.2408.16604>
- García Bustos, J. E., Baeza, C., Schiele, B. B., Rivera, V., Masserano, B., Orchard, M. E., ... Perez, A. (2025, 2). A novel data-driven framework for driving range prognostics in electric vehicles. *Engineering Applications of Artificial Intelligence*, 142, 109925. doi: <https://doi.org/10.1016/j.engappai.2024.109925>
- Hou, J., Yang, M., Wang, D., & Zhang, J. (2020, February). Fundamentals and challenges of lithium ion batteries at temperatures between -40 and 60 °C. *Advanced Energy Materials*, 10(18). doi: <https://doi.org/10.1002/aenm.201904152>
- Hu, X., Xu, L., Lin, X., & Pecht, M. (2020, February). Battery lifetime prognostics. *Joule*, 4(2), 310–346. doi: <https://doi.org/10.1016/j.joule.2019.11.018>
- International Energy Agency. (2025). *Global EV outlook 2025*. Paris. Retrieved from <https://www.iea.org/reports/global-ev-outlook-2025> (Licence: CC BY 4.0)
- McLachlan, G. J., & Krishnan, T. (2008). *The em algorithm and extensions*. John Wiley & Sons.
- Oh, G., Leblanc, D. J., & Peng, H. (2022). Vehicle energy dataset (ved), a large-scale dataset for vehicle energy consumption research. *IEEE Transactions on Intelligent Transportation Systems*, 23(4), 3302–3312. doi: <https://doi.org/10.1109/TITS.2020.3035596>
- Pola, D. A., Navarrete, H. F., Orchard, M. E., Rabié, R. S., Cerda, M. A., Olivares, B. E., ... Pérez, A. (2015). Particle-Filtering-Based Discharge Time Prognosis for Lithium-Ion Batteries With a Statistical Characterization of Use Profiles. *IEEE Transactions on Reliability*, 64(2), 710–720. doi: <https://doi.org/10.1109/TR.2014.2385069>
- Pérez, A., Quintero, V., Rozas, H., Jaramillo, F., Moreno, R., & Orchard, M. (2017). Modelling the degradation process of lithium-ion batteries when operating at erratic state-of-charge swing ranges. In *2017 4th international conference on control, decision and information technologies (codit)* (p. 0860–0865). doi: <https://doi.org/10.1109/CoDIT.2017.8102703>
- Rahman, T., & Alharbi, T. (2024, June). Exploring lithium-ion battery degradation: A concise review of critical factors, impacts, data-driven degradation estimation techniques, and sustainable directions for energy storage systems. *Batteries*, 10(7), 220. doi: <https://doi.org/10.3390/batteries10070220>
- Reynolds, D. (2015). Gaussian mixture models. In *Encyclopedia of biometrics* (pp. 827–832). Springer.
- Spitthoff, L., Shearing, P. R., & Burheim, O. S. (2021). Temperature, ageing and thermal management of lithium-ion batteries. *Energies*, 14(5). doi: <https://doi.org/10.3390/en14051248>
- World Weather Online. (2025). *World weather online api documentation*. <https://www.worldweatheronline.com/developer/>. (Accessed: 25 Jun 2025)
- Yeon, H., Eom, T., Jang, K., & Yeo, J. (2023, March). Dtumos, digital twin for large-scale urban mobility operating system. *Scientific Reports*, 13(1). doi: <https://doi.org/10.1038/s41598-023-32326-9>
- Zhou, C., Qian, K., Allan, M., & Zhou, W. (2011). Modeling of the cost of ev battery wear due to v2g application in power systems. *IEEE Transactions on Energy Conversion*, 26(4), 1041–1050. doi: <https://doi.org/10.1109/TEC.2011.2159977>