

Risk-Aware Optimization of Charging Time and Route Selection for Electric Vehicles Under Uncertainty

Jorge E. García Bustos¹, Bruno Masserano², Ricardo Salas Espiñeira³, Benjamín Brito Schiele⁴, Leonardo Baldo⁵, Vicente Pinochet⁶, Francisco Jaramillo-Montoya⁷, Heraldo Rozas⁸, Aramis Perez⁹, and Marcos E. Orchard¹⁰

^{1,2,3,6,7,8,10} *Department of Electrical Engineering,
Faculty of Physical and Mathematical Sciences, University of Chile, Santiago, Chile*

*jorgegarcia@ug.uchile.cl
bruno.masserano@ug.uchile.cl
ricardo.salas.e@ug.uchile.cl
vicente.pinochet.r@ug.uchile.cl
francisco.jaramillo@uchile.cl
heraldo.rozas@ug.uchile.cl
morchard@ing.uchile.cl*

⁴ *Intelligent System Prognostics Group, Aerospace Structures and Materials Department,
Faculty of Aerospace Engineering, Delft University of Technology, Delft, 2629HS, The Netherlands
bbritoschiele@tudelft.nl*

⁵ *Department of Mechanical and Aerospace Engineering,
Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Torino, Italy
leonardo.baldo@polito.it*

⁹ *School of Electrical Engineering,
University of Costa Rica, San José, Costa Rica
aramis.perez@ucr.ac.cr*

ABSTRACT

Effective decision-making during fast-charging sessions is becoming increasingly critical as Electric Vehicles (EVs) are deployed at scale. Drivers must make operational decisions under uncertainty arising from charging-time requirements, travel constraints, and the risk of battery depletion before destination. Moreover, from the EV driver's point of view, the reality of popular EV charging stations is far from ideal: charging station billing schemes, charging power decreases over time, and the applicable charging protocol depends on both the charger and the vehicle. In this setting, a driver charging at a given station who intends to reach a predefined destination faces two key operational questions: how long to charge the vehicle and which route to take, given that multiple feasible routes may be available. Accordingly, the driver must determine the required charging time to reach

the destination, accounting for uncertainty in the estimated State of Charge (SoC) and stochasticity in route-dependent energy demand, while ensuring that the route can be completed without violating a voltage-based feasibility condition associated with power cut-off. Against this backdrop, we formulate the charging decision as a probabilistic optimization problem that captures the trade-off between charging time, travel time, and the probability of energy shortfall under a user-defined level of risk tolerance. The problem is centered on estimating the likelihood of an End-of-Power-Availability (EPA) event, thereby enabling route-aware range prediction under uncertainty. We conduct experiments using real charging curves, time-based tariffs, and two alternative routes in Costa Rica. Results show a clear trade-off between charging duration and EPA probability and consistent improvements over simple heuristics such as charging to full capacity or targeting a fixed SoC threshold. These results position the proposed approach as a practical decision-support tool for EV users, enabling charging and routing decisions to be made explicitly under uncertainty and user-defined risk preferences.

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.

1. INTRODUCTION

The rising adoption of Electric Vehicles (EVs) worldwide is reshaping the transportation sector, not only redefining how people travel but also influencing urban mobility patterns, the planning and operation of charging infrastructure, and institutional global policy agendas (Zhan, Liao, Deng, Wang, & Yeh, 2025; Unterluggauer, Rich, Andersen, & Hashemi, 2022). In this sense, EVs should be seen as a pillar of a broader transition to shift greenhouse gas emissions away from urban centers, to deep decarbonize the energy system, and to develop smarter, more flexible utilities (R. Liu et al., 2024).

This paradigm shift requires a rapidly evolving infrastructure, without which a large-scale transition in mobility cannot be effectively enabled (Noel, de Rubens, Kester, & Sovacool, 2020; Biresselioglu, Kaplan, & Yilmaz, 2018). Such infrastructure includes coordinated financial transition plans, targeted investments in critical sectors, and a widespread charging network that supports users in shifting from an internal-combustion paradigm to a battery-powered one (Unterluggauer et al., 2022). This transition remains challenging, as range anxiety and charging uncertainty (e.g., availability, waiting times, and charging-point reliability) continue to be major barriers to EV adoption (Egbue & Long, 2012). As a result, public charging stations are becoming an increasingly visible component of modern mobility, although behind the simple act of plugging in lies a complex operational ecosystem that must be managed for the EV transition to scale efficiently and seamlessly. The International Energy Agency (IEA) reports that Europe’s public charging network grew by more than 35% in 2024 relative to 2023, surpassing 1 million charging points (International Energy Agency (IEA), 2025). At the same time, the 2025 European EV Charging Report estimates that Europe will need around 8.8 million chargers by 2030, highlighting a substantial gap that must be closed through faster deployment, improved operational efficiency, or both (gridX, 2025). The same report also reports user behavior consistent with congestion-aware decision-making: drivers strongly prefer home charging for convenience, whereas speed becomes critical when public charging is needed, particularly on long trips, for which around 70% of users plan charging stops in advance. These trends suggest that decisions about where to charge, how long to charge, and which route to take are inherently interconnected and must account for charging requirements, charging protocols, travel constraints, and the risk of battery depletion before reaching the destination.

In particular, this paper addresses a compelling research and practical question: given an EV that is already plugged in at a charging station, how should the driver jointly decide how long to continue charging (i.e., the target state-of-charge at departure) and which route to take from that station to the

final destination, so as to balance travel time, expected waiting/charging time, energy feasibility, and overall trip cost under uncertainty (Hoen, Díez-Gutiérrez, Babri, Hess, & Tørset, 2023)? From a broader perspective, this problem can be interpreted as a driver-centric decision-making task under uncertainty, in which operational choices must be made by balancing time efficiency, energy feasibility, and the tolerance for mission failure risk.

This decision-making problem under uncertainty involves several intertwined factors that are critical for informed and effective choices. In this operational setting, the driver must account for charging-time requirements, travel constraints, the risk of battery depletion before reaching the destination, uncertainty in the estimated State of Charge (SoC), station-specific pricing schemes, battery health, the nonlinear nature of the charging process (e.g., charging-power tapering), and protocol-dependent interactions between charger and vehicle. Under these conditions, the task can be naturally formulated as an optimization problem in which charging duration and route selection must be jointly determined under uncertainty, while explicitly reflecting the user’s risk attitude.

This decision becomes particularly non-trivial in modern fast-charging environments, which are playing an increasingly prominent role in public EV charging infrastructure. Recent reports (Brown, Cappellucci, Gaus, & Buleje, 2024) show that a growing share of newly deployed public chargers corresponds to fast or ultra-fast technologies, consistent with users’ preference for reducing waiting times during long-distance or time-critical trips (Gnann et al., 2018; Hanig et al., 2025). In these settings, charging stations often apply time-based tariffs, billing users per minute rather than per unit of energy, as in Costa Rica’s public fast-charging network (Autoridad Reguladora de los Servicios Públicos (ARESEP), 2023). As a result, user cost depends directly on charging duration rather than on the effective energy delivered.

At the same time, charging power is not constant throughout the session. Due to battery electrochemical limits, thermal constraints, and protocol-specific control strategies, charging power typically decreases as the SoC increases (Tomaszewska et al., 2019; Clar-Garcia, Fabra-Rodriguez, Campello-Vicente, & Velasco-Sanchez, 2025). This results in a nonlinear and vehicle-dependent relationship between charging time and energy stored, shaped jointly by the charger’s capabilities and the vehicle’s charging protocol (Löbel, Borndörfer, & Weider, 2023). As a consequence, additional waiting time yields diminishing marginal energy gains, tightly coupling charging duration, cost, and achievable driving range. Under such conditions, simple heuristics, such as charging to full capacity or departing at a fixed SoC threshold, can lead to inefficient decisions, excessive waiting times, increased congestion at charging stations, and a suboptimal user experience.

Uncertainty further complicates this planning task. The SoC reported by the Battery Management System (BMS) is inherently uncertain because it depends on estimation algorithms that are affected by temperature, aging, and recent operating history. Likewise, the energy required to complete a trip depends on traffic conditions, road geometry, driving behavior, and environmental factors. When multiple candidate routes connect the charging station to the destination, each route is characterized by distinct travel times and energy-demand distributions. The driver, therefore, faces a fundamental trade-off among charging duration, travel time, and robustness to adverse energy realizations (Hoen et al., 2023).

In this work, this trade-off is formulated as a multi-objective optimization problem that balances waiting and travel times against the risk of insufficient energy to complete the trip. Rather than imposing energy feasibility as a deterministic constraint, the framework adopts a probabilistic perspective, using events such as End-of-Power-Availability (EPA) or violations of the Maximum Driving Range (MDR) (J. E. Bustos et al., 2025; J. E. G. Bustos et al., 2026) as interpretable measures of mission-failure risk. These risk measures provide a concrete representation of uncertainty and user-defined risk tolerance. The resulting solutions define a Pareto frontier, where lower risk comes at the cost of longer charging times, and vice versa. By explicitly modeling this frontier, the proposed approach supports informed, user-adaptive decisions that reflect realistic fast-charging behavior, route-dependent energy uncertainty, and individual preferences regarding time efficiency and reliability.

In conclusion, this paper makes three main contributions: (i) it proposes a unified decision-making and optimization framework that captures realistic fast-charging conditions to jointly determine charging duration and route selection under uncertainty; (ii) it introduces a risk-aware multi-objective formulation that balances charging time, travel time, and the probability of EPA; and (iii) it provides an experimental validation using real charging curves and two candidate routes to benchmark the proposed strategy. Beyond route-aware range prediction, this work contributes an interpretable decision-making framework to support EV users in realistic fast-charging scenarios.

2. THEORETICAL BACKGROUND

2.1. Battery Model

To determine when an EV approaches the limit of its available autonomy, the future evolution of battery voltage must be estimated under the expected operating conditions. In practical Battery Management Systems (BMSs), the minimum allowable terminal voltage is typically used as the main criterion for defining the end of usable capacity, since once this safety threshold is reached, further discharge is prevented to avoid cell damage. Therefore, predicting the battery voltage trajec-

tory is essential for assessing whether the vehicle can complete a planned trip segment or requires recharging.

To characterize the dynamic behavior and low-frequency voltage response of the EV lithium-ion battery, this framework employs the Thévenin equivalent circuit presented in Burgos-Mellado, Orchard, Kazerani, Cárdenas, and Sáez (2016). The model represents the polarization voltage drop primarily through the electrolyte's ohmic resistance, omitting capacitive effects due to the assumed sampling frequency.

The system is defined by a discrete-time state-space representation that depends on the battery's SoC. The state transition equation describes the evolution of SoC over time:

$$SoC_k = SoC_{k-1} - \frac{I_{k-1} \cdot \Delta t_k}{Q_{max}} + \omega_k, \quad (1)$$

where I_{k-1} is the electrical current, Δt_k is the sampling period, Q_{max} is the maximum capacity of the battery, and ω_k represents process noise accounting for unmodeled dynamics such as hysteresis.

The measurement equation maps the internal state to the observable terminal voltage (V_k). It incorporates the open-circuit voltage (V_{oc}), the voltage drop across the internal resistance (R_{int_k}), and measurement noise (ϵ_k):

$$V_k = V_{oc}(SoC_k) - I_k \cdot R_{int_k}(SoC_k, I_k) + \epsilon_k. \quad (2)$$

The internal resistance R_{int_k} is a nonlinear function of both the current SoC and the electric current. The instantaneous electric power (P_k) at the battery terminals is computed as:

$$P_k = V_{oc}(SoC_k) \cdot I_k - R_{int_k} \cdot I_k^2. \quad (3)$$

The open-circuit voltage V_{oc} exhibits a highly nonlinear relationship with the SoC, strongly shaping discharge behavior prior to an EPA event. This relationship is parameterized using an empirical function:

$$V_{oc}(SoC_k) = \bar{\theta}_1 + (\bar{\theta}_2 - \bar{\theta}_1) \cdot e^{\bar{\theta}_5 (SoC_k - 1)} + \bar{\theta}_3 \cdot \bar{\theta}_1 (SoC_k - 1) + (1 - \bar{\theta}_3) \bar{\theta}_1 \left(e^{-\bar{\theta}_4} - e^{-\bar{\theta}_4 \sqrt{SoC_k}} \right), \quad (4)$$

where parameters $\bar{\theta}_1, \bar{\theta}_2, \dots, \bar{\theta}_5$ are derived offline through optimization techniques applied to experimental discharge data (Zhang & Chow, 2010). During operation, establishing a minimum threshold for the terminal voltage constraint directly enables the probabilistic calculation of maximum range feasibility.

2.2. Machine Learning Architectures

The future battery voltage trajectory depends directly on the power demanded by the EV during operation, which in turn is influenced by route characteristics, vehicle dynamics, and the

environmental conditions encountered at a given time. Estimating future battery power demand is therefore a necessary step toward predicting the voltage trajectory. Since these operating conditions evolve sequentially along a trip and exhibit strong temporal dependencies, data-driven models capable of capturing time-series dynamics are required. In this work, neural network architectures are employed to model the EV's future power demand by leveraging the sequential nature of the driving conditions. The main architecture adopted in this study is presented next.

Recurrent Neural Networks (RNNs) are widely used in PHM to model multivariate degradation signals with temporal dependencies. In a standard RNN, the hidden state is recursively updated as new time steps are processed, allowing past information to influence subsequent predictions. However, long input sequences often lead to vanishing or unstable gradients, limiting the ability of simple RNNs to learn long-term degradation patterns. For this reason, Gated Recurrent Units (GRUs) are commonly adopted in prognostics. GRUs introduce update and reset gates that regulate the retention and removal of past information, improving the modeling of long-term temporal dependencies while using fewer parameters than LSTM-based architectures. This makes them suitable for RUL estimation and health-state modeling from sequential condition-monitoring data (Chen, Jing, Chang, & Liu, 2019; Baptista, Mishra, Henriques, & Prendinger, 2024).

Regression trees provide a complementary approach for PHM tasks based on engineered features or health indicators. They recursively partition the input space through binary decision rules. Predictions are then obtained from the statistics of the samples contained in the reached leaf. Moreover, practical implementations often rely on heuristic procedures such as recursive binary splitting and pruning (Sohil, Sohali, & Shabbir, 2022). In this context, gradient boosting improves this framework by training trees sequentially, so that each new tree reduces the residual error of the current ensemble. The Light Gradient Boosting Machine (LightGBM) is often used as an efficient implementation of these gradient-boosted regression trees. LightGBM accelerates training through Gradient-based One-Side Sampling, which prioritizes samples with large gradients during split estimation, combining mutually exclusive features to reduce the effective dimensionality of the problem (Ke et al., 2017). Such properties are useful in PHM applications where RUL models may rely on multiple sensor-derived features or health indicators (H. Liu et al., 2022; Ke et al., 2017).

The selection of GRU and LGBM in this work is not intended to introduce new machine-learning architectures, but rather to use models that are consistent with the structure of the prediction problem and with the computational requirements of the proposed decision framework. GRUs are adopted because the driving process evolves sequentially along the

route, and therefore the speed and demand prediction stages must capture temporal dependencies without adding unnecessary model complexity. Compared with LSTM-based alternatives, GRUs provide a lighter recurrent structure while still retaining the ability to model relevant short- and medium-term dependencies in the route profile. LGBM is used as a complementary model because part of the power-demand estimation problem is naturally represented through tabular route and operational features, for which gradient-boosted decision trees offer a favorable trade-off between predictive performance, robustness, and computational efficiency. In this sense, both models were selected as practical components of a route-aware prognostic pipeline, rather than as isolated forecasting tools.

2.3. Maximum Driving Range

Building upon the route-demand prediction models described in the previous section, the MDR methodology integrates data-driven demand prediction with a physics-based battery model to estimate the likelihood of an EV reaching its autonomy limit along a route. The MDR framework predicts the segment of a route where an EV will likely experience an EPA event. To achieve this, the driving route is divided into sequential segments, and a GRU network together with an LGBM model are used to generate speed and power consumption trajectories (J. E. Bustos et al., 2025). These predicted trajectories are then used as inputs to the battery model to compute the voltage evolution and evaluate the probability of an EPA event for each route segment.

The original MDR from J. E. Bustos et al. (2025) relies on Monte Carlo (MC) simulations to evaluate the probability of critical battery conditions. The probability of an EPA event at segment k , denoted as $\mathbb{P}(EPA_k|\mathcal{R})$, is calculated by evaluating the event where terminal voltage V_k falls below a defined voltage threshold $V_{threshold}$ over the space of power trajectories, sensor noise, and process noise:

$$\begin{aligned} \mathbb{P}(EPA_k|\mathcal{R}) &= \int_{\mathcal{P}_{1:k} \times \mathcal{H} \times \Omega_{1:k}} \mathbb{P}(V_k \leq V_{threshold} | P_{1:k}, \eta_k, \omega_{1:k}) \cdot \dots \\ &\quad \dots \prod_{i=1}^{k-1} (1 - \mathbb{P}(V_i \leq V_{threshold} | P_{1:i}, \eta_i, \omega_{1:i})) \cdot \dots \\ &\quad \dots \mathbb{P}(P_{1:k}, \eta_{1:k}, \omega_{1:k} | \mathcal{R}) d(P_{1:k}, \eta, \omega_{1:k}). \end{aligned} \quad (5)$$

This numerical approximation requires extensive sampling of all stochastic variables. The reliance on extensive sampling imposes a substantial computational cost that restricts its utility to offline scenarios.

The Uncertain Event Likelihood Function (UELF) framework introduced in J. E. G. Bustos et al. (2026) replaces MC

simulations by analytically incorporating additive uncertainty in sensor noise. Instead of using a fixed deterministic threshold, UELF defines event likelihoods by expressing the probability of crossing the threshold via the cumulative distribution function of a Gaussian distribution (Acuña-Ureta & Orchard, 2023):

$$\begin{aligned} \mathbb{P}(V_k \leq V_{threshold} | P_{1:k}, \eta_k, \omega_{1:k}) \\ = 1 - \phi_{\mu_k, \sigma_\eta}(v_{oc}(SoC_k)), \end{aligned} \quad (6)$$

where $\phi_{\mu_k, \sigma_\eta}(\cdot)$ represents the cumulative distribution function with a mean of

$$\mu_k = V_{threshold} + \sqrt{P_k \cdot R}, \quad (7)$$

and a variance equal to the sensor noise variance σ_η^2 .

By applying the UELF formulation, the computation marginalizes over the SoC and power, eliminating the need to explicitly sample sensor noise. The final probability calculation is given by:

$$\begin{aligned} \mathbb{P}(EPA_k | SoC_0, \mathcal{R}) \\ = \int_{\mathcal{P}_{1:k} \times \Omega_{1:k}} (1 - \phi_{\mu_k, \sigma_\eta}(v_{oc}(SoC_k))) \cdots \\ \cdots \prod_{i=1}^{k-1} \phi_{\mu_i, \sigma_\eta}(v_{oc}(SoC_i)) \cdots \\ \cdots \mathbb{P}(P_{1:k}, \omega_{1:k} | \mathcal{R}) d(P_{1:k}, \omega_{1:k}). \end{aligned} \quad (8)$$

This abstraction reduces algorithmic complexity and computational overhead while guaranteeing asymptotic stochastic convergence, enabling the real-time execution of the MDR-UELF algorithm (J. E. G. Bustos et al., 2026).

2.4. Chance-Constrained Optimization

The MDR methodology provides a probabilistic estimate of the risk of an EPA event along a route, given a specific initial SoC. In a route planning context, this risk can be controlled through the selection of the initial SoC before departure, since when the vehicle is at a charging station, the driver may decide the charge level at which the trip will start. In this context, it becomes necessary to determine the minimum initial SoC that keeps the risk of an EPA event below a predefined risk-aversion threshold.

Because the predicted risk depends on uncertain variables such as traffic conditions, driving behavior, and environmental factors, this problem cannot be formulated using purely deterministic constraints. More generally, when the parameters governing a system are uncertain, deterministic constraints of the form $g(x, \xi) \leq 0$, where ξ is a random variable, cannot be enforced in the classical sense. Chance-constrained programming (Nemirovski & Shapiro, 2007) addresses this issue through a probabilistic reformulation that

requires the constraint to hold with at least a prescribed probability $\alpha \in [0, 1]$:

$$\mathbb{P}(g(x, \xi) \leq 0) \geq \alpha, \quad (9)$$

where the parameter α directly encodes the decision-maker's risk aversion: higher values of α correspond to more conservative feasibility requirements, at the cost of a more restricted feasible set. This probabilistic reformulation is particularly well-suited to EV range planning, where energy consumption along a route is inherently stochastic due to traffic variability, driver behavior, and environmental conditions.

In the specific EV range-planning problem addressed in this work, the probabilistic feasibility condition is implemented through the complementary failure event, namely the probability of an End-of-Power-Availability event. Therefore, instead of directly enforcing a success probability α , the optimization constrains the EPA risk to remain below an admissible tolerance δ , with $\delta = 1 - \alpha$ when the success and failure events are defined as complementary route-level events. In this notation, lower values of δ represent more conservative decisions, since they allow a smaller probability of voltage-threshold violation along the route.

In the case of this work, a central property that makes chance constraints tractable is the monotonicity of the constraint satisfaction probability with respect to the initial SoC level. Intuitively, a higher departure SoC shifts the battery voltage distribution upward throughout the route, so the probability of satisfying a minimum voltage requirement at the destination is non-decreasing in SoC_{dep} . This monotonic structure implies that the feasible set in the SoC_{dep} space is an interval of the form $[SoC_{dep}^*, 1]$, where SoC_{dep}^* is the unique minimum departure charge that satisfies the EPA risk constraint at the prescribed tolerance δ .

The bisection method is a numerical technique used to locate the root of a continuous function within a bounded interval. The method is based on the intermediate value theorem, which states that if a continuous function $f(x)$ takes values of opposite sign at two points a and b , then there exists at least one root x^* in the interval $[a, b]$ such that $f(x^*) = 0$.

Given an initial interval $[a, b]$ where $f(a)$ and $f(b)$ have opposite signs, the method proceeds by iteratively halving the interval. At each iteration, the midpoint is computed as

$$x_m = \frac{a + b}{2}. \quad (10)$$

The function is then evaluated at the midpoint. If $f(x_m)$ has the same sign as $f(a)$, the root must lie in the interval $[x_m, b]$; otherwise, it lies in $[a, x_m]$. The interval is therefore reduced by selecting the subinterval that preserves the sign change condition. This process is repeated until the interval width or the function value satisfies a predefined tolerance.

The bisection method is particularly suitable for the charging optimization problem addressed in this work because the decision variable is one-dimensional and corresponds to the departure SoC. Moreover, the route feasibility probability is monotonic with respect to this variable: increasing the departure SoC shifts the predicted voltage trajectory upward and reduces the probability of violating the minimum voltage constraint. Therefore, the problem can be reformulated as finding the minimum SoC that satisfies the prescribed probabilistic feasibility condition within a bounded interval. Under this structure, bisection provides a robust and transparent solution strategy, since it does not require gradient information, is insensitive to the scale of the objective function, and guarantees convergence once the feasible interval is properly bracketed. Other alternatives, such as exhaustive grid search, evolutionary optimization, or gradient-based methods, could also be applied. However, a grid search would require a predefined resolution and may become inefficient for tighter tolerances; evolutionary methods would add unnecessary computational cost for a one-dimensional monotonic problem; and gradient-based methods are unnecessary because the feasibility boundary can be directly located through interval reduction. For this reason, bisection was selected as the most appropriate strategy for this specific chance-constrained charging problem.

3. METHODOLOGY

This work considers the decision-making problem of an EV driver who, while located at a charging station, must determine the optimal charging time required to reach a target destination while minimizing both the total travel time and the cost associated with the charging process, subject to the charging protocol imposed by the EV–charger pair (e.g., admissible power levels, time-varying charging profiles, and operational limits) and to a route-feasibility requirement that prevents the battery voltage from crossing a prescribed safety threshold during the trip. Accordingly, the proposed methodology should be understood as a decision-making architecture in which the charging target and the route choice are treated as coupled operational decisions rather than as independent planning stages. In this sense, the charging decision is not only an economic–time trade-off, but must also guarantee that the selected route can be completed without violating a minimum voltage constraint that reflects the driver’s safety margin and risk tolerance, i.e., that the route is feasible under the initial conditions the EV starts the mission.

An overview of the proposed methodology is illustrated in Figure 1. The framework starts from the trip origin and destination, from which a set of candidate routes is generated. Each route is decomposed into segments and processed through a feature extraction and segmentation pipeline. Then, the MDR-based probabilistic model is used to evaluate route feasibility under uncertainty, allowing the computation of the probability of satisfying a minimum voltage constraint along

the route. Based on this information, a chance-constrained optimization determines the minimum charging time required for each route to meet the user-defined risk tolerance. Finally, the decision-making layer selects the route that minimizes the total mission time, while ensuring that the associated charging time guarantees route feasibility.

Let $\mathcal{R} = \{r_1, r_2, \dots, r_{N_r}\}$ denote the set of feasible routes that connect the charging station to the destination. Each route $r \in \mathcal{R}$ is composed of an ordered sequence of segments

$$r = \{s_1, s_2, \dots, s_{N_s(r)}\}, \quad (11)$$

where $N_s(r)$ is the number of segments of route r . The energy and power demands and driving conditions may vary across segments and routes.

The core problem addressed in this work is the determination of the charging time t_c that balances charging cost and travel performance under uncertainty, while explicitly accounting for the driver’s risk preferences. In this work, the driver’s risk preferences refer specifically to the degree of conservatism adopted when managing energy uncertainty along the route, which is encoded through an admissible EPA risk tolerance and safety margins on battery constraints. In particular, the charging decision must ensure, under a prescribed EPA risk tolerance, that the vehicle reaches the destination with a user-defined terminal voltage requirement, which acts as a proxy for the desired residual energy margin at arrival. Since the EV power cut-off is ultimately triggered by a voltage disconnection, the feasibility condition is naturally posed in terms of voltage. The driver specifies a desired energy reserve at the destination, and this requirement is mapped to an admissible terminal voltage threshold consistent with that reserve. Accordingly, the confidence bands of the autonomy forecast are required to cover at least the voltage level associated with the user’s target SoC at the destination. This probabilistic formulation allows the driver to trade off charging cost and travel efficiency against the risk of insufficient energy availability along the selected route.

3.1. Route Feasibility Evaluation

The proposed methodology relies on the MDR framework to perform a probabilistic prediction of vehicle autonomy. Because feasibility is ultimately enforced through a minimum voltage requirement, the charging decision must be supported by a risk measure that is both operational and interpretable from the user’s perspective. Quantifying risk in terms of battery voltage provides a direct link between the model output and the decision constraint, transforming the autonomy prediction into actionable information about how conservative the plan is and how close it operates to the cut-off condition. In this way, the risk metric does not only reflect uncertainty, but also adds an interpretability layer that helps the driver understand the safety margin implied by a given charging time

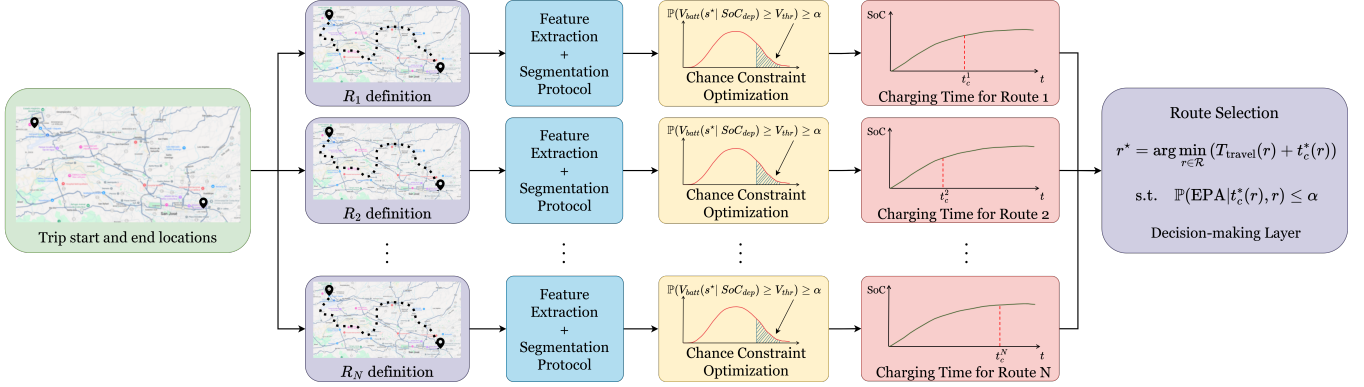


Figure 1. Graphical representation of the charging time optimization process. The optimal charging time is determined as the minimum charging duration that ensures route feasibility at the destination under the prescribed risk tolerance.

and route choice. For a given route r and segment s , the MDR framework characterizes the probability density function (PDF) of the battery voltage,

$$p_V(v | s, r), \quad (12)$$

which captures the uncertainty associated with multiple sources that affect the voltage response along the route. In particular, part of the uncertainty is linked to the battery model and its parameters, part is associated with voltage sensing and measurement noise, and part emerges from the variability embedded in the input features used by the speed and energy models. While these contributions are not explicitly separated, they are jointly propagated through the predictive pipeline and reflected in the voltage distribution through the stochastic dropout mechanism.

Let V_{th} denote a predefined voltage threshold associated with a minimum acceptable state of charge. This threshold plays a central role in the proposed decision-making framework because it is the mechanism through which the driver's risk aversion is made explicit and operational. This threshold is selected according to the driver's risk aversion and reflects the minimum voltage required to safely complete a route segment. A more conservative driver adopts a higher V_{th} , which increases the safety margin to the power cut-off condition and reduces the probability of disconnection, at the expense of requiring longer charging times.

This probability quantifies the likelihood that the vehicle can traverse segment s without violating the voltage constraint, thereby providing a probabilistic measure of range feasibility.

3.2. Charging Optimization Problem

The objective is to find the minimum charging time t_c that ensures that the route is feasible from the MDR framework standpoint, under a predefined admissible EPA risk tolerance δ and a minimum voltage threshold V_{thr} . Viewed from a decision-making standpoint, this step determines the mini-

mum charging action required to make a route admissible under uncertainty.

Since the charging protocol defines a one-to-one relation between charging time and departure state of charge, this problem can be equivalently formulated in terms of the departure state of charge SoC_{dep} reached after charging. Then, the optimization problem consists of finding the minimum departure state of charge that satisfies the feasibility constraint imposed by the MDR framework for a given route.

This is equivalent to:

$$\min_{SoC_{dep}} T_{travel}(r) + t_c(SoC_{dep}) \quad (13)$$

Here, given a departure state of charge SoC_{dep} , the MDR framework is used to propagate uncertainty along the route and compute the cumulative probability of surpassing the voltage threshold at the destination segment s^* . The corresponding feasibility constraint is expressed as:

$$\text{s.t. } P_{EPA}(s^* | SoC_{dep}) \leq \delta \quad (14)$$

where $P_{EPA}(s^* | SoC_{dep})$ denotes the cumulative probability of reaching an EPA event up to the destination segment s^* , and δ is the maximum admissible EPA risk defined by the user. Equivalently, this condition requires the probability of completing the route without violating the voltage constraint to be at least $1 - \delta$. Due to the monotonic relation between SoC_{dep} and the cumulative EPA risk obtained from the MDR framework, the optimization is solved using a bisection search.

The process starts by evaluating the current state of charge without additional charging, SoC_0 . If the route is already feasible at SoC_0 , then no charging time is required, and the solution is obtained directly.

If the route is not feasible at SoC_0 , the maximum admissible state of charge, $SoC = 1$, is evaluated. If the route remains

Algorithm 1 Bisection update of departure SoC based on terminal EPA risk

```

1: Initialize  $SoC_{low} = SoC_0$ ,  $SoC_{high} = 1$ 
2:
3: while  $SoC_{high} - SoC_{low} > \varepsilon$  do
4:
5:    $SoC_{mid} = (SoC_{low} + SoC_{high})/2$ 
6:    $p_{EPA} = Risk(MDR(s^* | SoC_{mid}))$ 
7:
8:   if  $p_{EPA} \leq \delta$  then
9:      $SoC_{high} \leftarrow SoC_{mid}$   $\triangleright$  feasible, reduce SoC
10:  else
11:     $SoC_{low} \leftarrow SoC_{mid}$   $\triangleright$  infeasible, increase SoC
12:  end if
13:
14: end while
15:
16:  $SoC_{dep}^{opt} = SoC_{high}$ 
    
```

infeasible even at full charge, the route is declared infeasible under the prescribed EPA risk tolerance. However, if it becomes feasible at full charge, then a feasible solution exists, but it may not be optimal in terms of minimum charging time.

The bisection procedure is then applied over the interval $[SoC_0, 1]$, as described in Algorithm 1. At each iteration, a midpoint departure SoC is tested using the MDR framework to evaluate the cumulative risk at the destination. If the midpoint is feasible, the upper bound is reduced; otherwise, the lower bound is increased. The interval is progressively reduced until the tolerance ε is satisfied, yielding the minimum feasible departure SoC.

3.3. Charging protocol definition

Once the optimal departure state of charge SoC_{dep}^{opt} has been determined, the corresponding charging time can be obtained by defining the charging protocol for the specific vehicle–charger pair. The complete charging function can be defined using the negotiated operating limits and the vehicle’s battery model. In practice, the charging trajectory $SoC = \tau(t_c)$ depends on the admissible power, current, and voltage levels that result from the interaction between the electric vehicle and the charging station.

Negotiation of operating constraints: During the charging handshake, both the vehicle and the charger communicate their maximum allowable limits. Then the negotiated constraints are defined as

$$I_{max} = \min(I_{max}^{ev}, I_{max}^{ch}), \quad (15)$$

$$V_{max} = \min(V_{max}^{ev}, V_{max}^{ch}), \quad (16)$$

$$P_{max} = \min(P_{max}^{ev}, P_{max}^{ch}). \quad (17)$$

These negotiated limits define the admissible operating region for the charging process.

Charging protocol construction: Based on the negotiated constraints and the selected charging protocol, the charging curve can be generated using the battery model. The battery terminal voltage is represented using a Thevenin equivalent model,

$$V_{term}(t) = V_{oc}(SoC(t)) + I(t)R_{int}, \quad (18)$$

where $V_{oc}(SoC)$ is the open-circuit voltage function and R_{int} is the internal resistance.

During the constant-current phase,

$$I(t) = I_{max}, \quad (19)$$

until the voltage constraint

$$V_{term}(t) = V_{max} \quad (20)$$

is reached. Afterward, the protocol transitions to a constant-voltage phase and the current decreases while keeping the $V_{term}(t)$ constant.

The optimal departure state of charge $SoC_{dep}^*(r)$ is then mapped to the optimal charging time through the charging protocol relation

$$t_c^{opt} = \tau^{-1}(SoC_{dep}^{opt}). \quad (21)$$

3.4. Route Selection

For each route $r \in \mathcal{R}$, the optimization problem described in the previous section yields an optimal charging time $t_c^*(r)$ that minimizes the charging cost while satisfying the probabilistic reachability constraint associated with the destination segment. Consequently, only routes for which a feasible charging solution exists under the prescribed risk level are considered admissible.

Among these admissible routes, the final routing decision is obtained by explicitly using the previously computed optimal charging times. The selected route is defined as the one that minimizes the total trip duration, given by the sum of driving time and optimal charging time.

$$r^* = \arg \min_{r \in \mathcal{R}} (T_{travel}(r) + t_c^{opt}(r)). \quad (22)$$

This formulation highlights that the routing decision is not independent of the charging problem, but rather directly informed by the risk-aware charging optimization. As a result, the proposed approach enables a coherent joint optimization of charging and routing decisions under uncertainty, accounting simultaneously for economic cost, travel time, and risk tolerance. Therefore, the final output of the methodology is not merely a feasible route, but a decision recommendation that jointly accounts for charging effort, travel duration, and mission-level risk.

4. RESULTS

The objective of this section is to evaluate the performance of the proposed methodology when applied to a real-world case study. In particular, the analysis aims to validate four key aspects of the framework. First, it seeks to verify that the methodology provides consistent and realistic estimates of the required charging time given the charging protocol and the probabilistic feasibility constraint. Second, it evaluates whether the predicted travel times for the candidate routes are consistent with the route characteristics and driving conditions. Third, the results assess whether the risk levels computed by the framework are consistent with the probabilistic reachability indicators provided by the MDR methodology. Finally, the study verifies that the route selected by the optimization procedure corresponds to the best alternative in terms of total trip duration, accounting for both travel time and optimal charging time. Taken together, these four aspects allow the proposed framework to be assessed not only as a predictive tool but also as a decision-making mechanism for selecting operationally preferable EV actions under uncertainty.

4.1. Case Study Description

To carry out this evaluation, the methodology is applied to a real-world case study conducted in Costa Rica using a 2018 Nissan Leaf equipped with a recently replaced battery pack with a nominal capacity of 30 kWh. Two routes were considered for this study. In both cases, they share the first 40 km and nearly the last 20 km. The only difference lies in the middle section of the routes. Although they are very similar, there is an important characteristic in this detour. The starting point is the University of Costa Rica campus, located in a valley in the middle of the country at an altitude of approximately 1,200 meters above sea level. The endpoint of the trips is Puerto Caldera, in the Central Pacific region of the country, at sea level. In the case of Route #1, it consists of a highway route where it is possible to drive at a constant speed without stopping for most of the route. In the case of Route #2, it is currently considered a secondary route, since it is a mountain road with many small ascents and descents, as well as curves, which limit the speed of the vehicle, making it a very demanding path for all types of vehicles. It is noteworthy that Route #2 has, in its mountainous section, a net ascent of about 400 meters, followed by a descent that occurs first along the mountain road, and after this, both routes meet at the junction. In this regard, both routes enable a controlled comparison in which performance differences arise from route geometry, operating conditions, and the resulting route-dependent energy demand and travel dynamics. In particular, the selected routes exhibit distinct topographical characteristics, as shown in Figure 3. Despite these differences, both routes have comparable lengths, ranging between approximately 90 and 95 km, as seen in Figure 2.

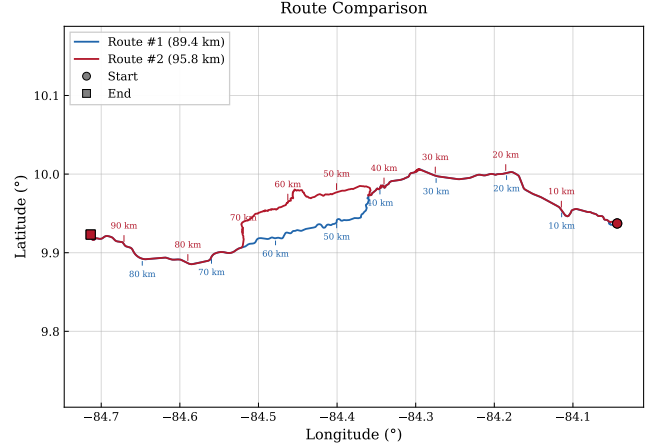


Figure 2. Geographical comparison of the two candidate routes used in the Costa Rica case study. Route #1 and Route #2 share the same origin and destination but differ in their path geometry and expected energy-demand conditions.

4.2. Battery Characterization

The battery model for the Nissan Leaf was trained using telemetry data collected in the Costa Rica metropolitan area. The dataset spans the period from September 2025 to March 2026 and comprises approximately 128,000 raw timestamped records. After preprocessing, a total of 107,216 valid samples remained, distributed across 488 recorded trips. Each sample contains the battery terminal voltage V (ranging from 340.0 to 400 V), the SoC ξ (from 5.6 % to 100 %), and an instantaneous current estimate \hat{I} derived from the cumulative driving and idling energy channels via

$$\hat{I} = \frac{\Delta E_{in,drv} + \Delta E_{in,idl} - \Delta E_{out,drv}}{\Delta t \cdot V}, \quad (23)$$

where all energy increments are expressed in kilowatt-hours and Δt in hours. The dataset was partitioned by trip index into training (70 %), validation (15 %), and test (15 %) subsets. The open-circuit voltage (OCV) model and internal resistance of the Nissan Leaf battery pack were identified using the methodology established in (Perez et al., 2021), which models the OCV as a nonlinear, monotonic function of the SoC.

The parameters R_{int} , $\bar{\theta}_1$, $\bar{\theta}_2$, $\bar{\theta}_3$, $\bar{\theta}_4$ and $\bar{\theta}_5$ were estimated by minimizing the Mean Squared Error (MSE) between measured and predicted terminal voltages over the training set, subject to the constraint that $V_{oc}(\xi)$ remain monotonically non-decreasing with ξ . Particle Swarm Optimization (PSO) was employed for this purpose, with a swarm of $N_p = 100$ particles evolved over 200 iterations using inertia weight $\omega = 0.72$ and acceleration coefficients $c_1 = c_2 = 1.49$, consistent with the canonical PSO formulation (Zhang & Chow, 2010). The search space bounds were set to physi-

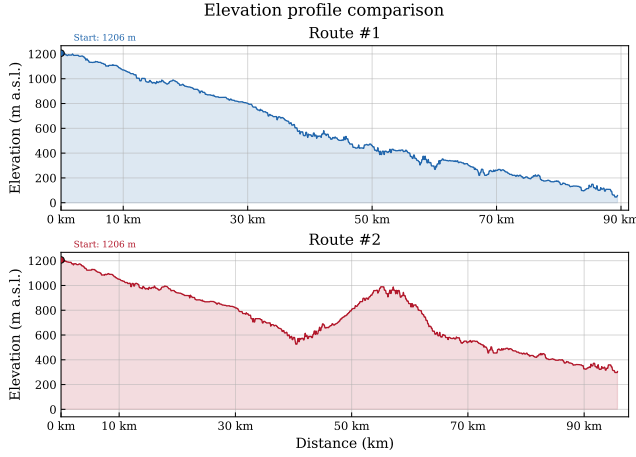


Figure 3. Elevation profile comparison for the two candidate routes. Route #2 presents stronger altitude variations than Route #1, which is expected to increase traction-energy variability and affect the risk-aware charging requirement.

cally meaningful ranges for the Nissan Leaf pack (nominal capacity ≈ 30 kWh, voltage range 325–400 V).

The PSO converged to a stable solution after approximately 50 iterations, achieving a best-cost value of 94.30 V^2 (RMSE $\approx 9.71 \text{ V}$ on the full dataset). The identified parameters are summarized in Table 1. The RMSE values disaggregated by split are 10.05 V (train), 8.11 V (validation), and 9.81 V (test), reflecting consistent generalization across unseen trips.

Table 1. Identified battery model parameters for the Costa Rica (Nissan Leaf) dataset via Particle Swarm Optimization.

Parameter	Symbol	Value	Unit
Threshold voltage	$V_{\text{threshold}}$	340.00	V
Internal resistance	R_{int}	0.005	Ω
Initial voltage	$\theta_1 (V_0)$	420.0000	V
Limit voltage	$\theta_2 (V_L)$	392.2745	V
Shape factor	$\theta_3 (\gamma)$	200.0000	—
Fitting coefficient	$\theta_4 (\alpha)$	0.1686	—
Fitting coefficient	$\theta_5 (\beta)$	25.0000	—

4.3. Charging protocol definition

With the battery model parameterized, the charging protocol associated with the EV–charger pair can be computed. This requires identifying the operational constraints imposed by both the vehicle and the charging station, as well as the charging strategy that governs their interaction (Huang, Soto, Ortiz, Arguello, & Perez, 2024). These limits include the maximum admissible voltage, current, and power levels supported by each system. The corresponding specifications are obtained from the vehicle and charger datasheets and are summarized in Table 2.

With the defined limits, along with the battery model and

Table 2. Charging constraints of the EV and charger in the case study; the resulting negotiated limits are emphasized in bold and correspond to the minimum of the pair.

	Voltage [V]	Current [A]	Power [kW]
Vehicle	400	130	47
Charger	500	100	20

the charging protocol, namely the constant-current/constant-voltage (CCCV) strategy, where the battery is first charged at a constant current until a voltage limit is reached, followed by a constant voltage phase with decreasing current, the complete charging trajectory can be reconstructed, and the charging time can be computed.

4.4. Route optimization

With the charging protocol defined, each route can be subsequently optimized. This process requires specifying a set of hyperparameters, including the pre-charging SoC, the battery voltage constraint and its associated uncertainty band, and the admissible EPA risk tolerance governing the maximum allowed probability of exceeding this limit. To ensure a fair and consistent comparison across routes, these parameters are kept identical across scenarios and summarized in Table 3.

Table 3. Parameters used for the optimization of the charging time for each route.

Parameter	Value
Max. bisections	50
Tolerance [h]	1.0×10^{-4}
Admissible EPA risk tolerance (δ)	0.5
Pre-charging SoC, SoC_0	0.55

Then, the proposed methodology is applied independently to each route to determine the minimum feasible departure SoC and the associated charging time at the prescribed EPA risk tolerance. The optimization yields an optimal departure SoC of 76.59% with a charging time of 0.287 [h] for Route #1, and an optimal departure SoC of 79.84% with a charging time of 0.317 [h] for Route #2.

The optimization results show that the two routes require different departure SoC levels in order to satisfy the same admissible EPA risk tolerance. This difference is explained by their route-dependent energy demand, induced by the distinct topographical profiles. In particular, Route #2 required a higher departure SoC than Route #1, which confirms that the proposed methodology is sensitive to route-specific operating conditions rather than relying on a fixed departure SoC rule. Although these optimized SoC values were obtained offline with high resolution, they could not be reproduced exactly during the real charging session. Therefore, both experimental runs departed from a common SoC value close to the

computed optimum. This slight discrepancy is representative of real operating conditions and reinforces the practical relevance of the validation, showing that the proposed methodology remains effective even when the executed charging condition only approximates the theoretical optimum.

Beyond validating the prediction accuracy, Figures 4 and 5 also provide a physical interpretation of the optimization outcome. Figure 4 shows that the SoC evolution along Route #1 follows a relatively smooth and gradual discharge pattern, with the measured trajectory remaining within the uncertainty envelope predicted by the MDR framework for essentially the entire route. By contrast, Figure 5 reveals a more demanding depletion process, characterized by a steeper and more irregular SoC evolution. This difference directly supports the higher optimized departure SoC obtained for Route #2. Taken together, these results show that the proposed framework is able to propagate route-dependent uncertainty while preserving consistency with the actual battery depletion observed in the field. Moreover, the contrast between Figures 4 and 5 is coherent with the distinct elevation profiles presented in Figure 3, and further supports the need for route-aware charging decisions. From a decision-making perspective, this is an important result because it shows that the charging target cannot be prescribed independently of the route, but must instead be adapted to the route-specific depletion dynamics anticipated by the framework.

Table 4 complements the visual assessment provided by Figures 4 and 5 with quantitative error metrics. In both routes, the SoC prediction errors remain low, with MAE and RMSE below 0.015 and 0.022, respectively. In addition, the uncertainty intervals achieve strong coverage of the measured trajectories, reaching 1.0000 for Route #1 and 0.9549 for Route #2. These results quantitatively support the consistency observed in Figure 4 for the smoother depletion case and in Figure 5 for the more demanding route, confirming that the proposed framework is able to track battery depletion with good accuracy while maintaining informative uncertainty bounds.

Table 4. Quantitative error metrics for the SoC prediction in the two candidate routes.

Metric	Route #1	Route #2
SoC MAE	0.0125	0.0148
SoC RMSE	0.0144	0.0216
Interval coverage	1.0000	0.9549
Final SoC abs. error	0.0156	0.0397

The route-level risk evolution is shown in Figure 6. The upper panel presents the cumulative EPA probability, whereas the lower panel highlights the marginal contribution of each route segment to the total mission-level risk. This distinction is important: while the cumulative curve indicates the




Figure 4. Predicted and ground-truth SoC evolution for Route #1 after charging. The shaded region denotes the MDR-based prediction interval, whereas the solid line represents the measured SoC trajectory.




Figure 5. Predicted and ground-truth SoC evolution for Route #2 after charging. The shaded region denotes the MDR-based prediction interval, whereas the solid line represents the measured SoC trajectory.

total probability of mission failure accumulated up to a given segment, the marginal curve reveals where along the route that risk is being generated. In both routes, the cumulative risk remains below the admissible threshold, confirming that the selected departure SoC values were sufficient to satisfy the imposed chance constraint. Moreover, Route #2 shows an earlier and more sustained activation of marginal risk, indicating that its risk accumulates over a longer fraction of the mission. In contrast, Route #1 concentrates risk later in the trip, with a sharper increase near the final segments. This result confirms that route feasibility is not only route-dependent, but also segment-dependent. In other words, the proposed framework does not merely classify a mission as feasible or infeasible, but also identifies where along the route the risk accumulates.

Table 5 summarizes the optimized charging and mission-level outcomes for both candidate routes. Route #1 is selected as the preferred alternative because it yields the lowest charging time, the lowest travel time, and therefore the minimum total mission time, while also presenting the lowest final EPA probability among the two feasible routes. This confirms that the proposed framework is able to support route selection through a joint evaluation of charging requirements, mission duration, and probabilistic feasibility. In other words, the selected solution should be interpreted as the output of

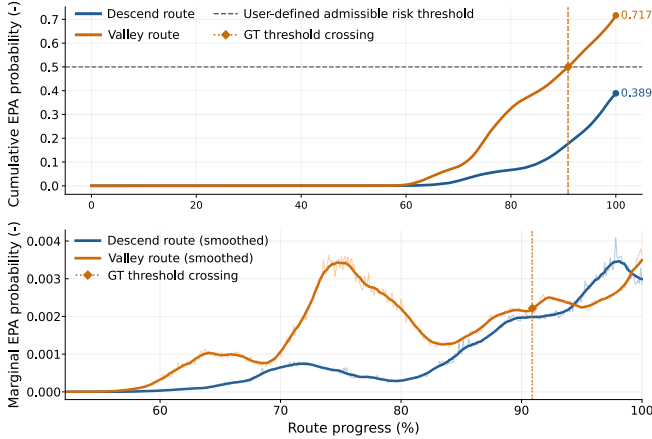


Figure 6. Comparison of the route-level MDR risk evolution for the two candidate routes after charging. The upper panel shows the cumulative EPA probability along route progress, while the lower panel shows the marginal EPA probability, highlighting the route segments where risk is locally activated. The dotted vertical lines indicate the ground-truth SoC threshold-crossing point, which serves as the reference for the onset of critical battery operation. The horizontal dashed line in the upper panel denotes the admissible EPA risk tolerance.

a structured decision process, in which charging and routing alternatives are screened according to a common probabilistic feasibility criterion and then compared through their total operational cost in time.

Table 5. Comparison of the optimized charging and mission-level outcomes for the two candidate routes.

Metric	Route #1	Route #2
Charging time [h]	0.287	0.317
Travel time [h]	1.999	2.103
Total mission time [h]	2.286	2.421
Energy charged [kWh]	5.950	6.350
Optimal SoC [%]	76.59	79.84
Final $P(\text{EPA})$	0.389	0.467
Selected route	Route #1	

4.5. Comparison with Simple Charging Heuristics

To further assess the practical value of the proposed strategy, Table 6 compares the optimized risk-aware decision against three simple charging heuristics: departing without additional charging, charging to a fixed SoC target of 85%, and charging to full capacity. The first two rows repeat the optimized results from Table 5 only as a reference point for the comparison. A strategy is considered feasible when its final EPA probability remains below the admissible EPA risk tolerance, i.e., $P(\text{EPA}) \leq \delta$.

The comparison shows that departing without additional

charging leads to the shortest mission time from a purely time-based perspective, but both routes become infeasible because the final EPA probabilities, 0.915 for Route #1 and 0.976 for Route #2, exceed the admissible EPA risk tolerance. This confirms that minimizing travel time alone is not sufficient when the probability of voltage-threshold violation is explicitly considered.

The fixed 85% SoC heuristic is feasible for both routes, but it does not adapt to route-specific requirements. For Route #1, the fixed target increases the total mission time from 2.286 h to 2.429 h, while only reducing the final EPA probability from 0.389 to 0.349. For Route #2, the same fixed target increases the total mission time from 2.421 h to 2.533 h, while reducing the final EPA probability from 0.467 to 0.439. Therefore, although the fixed-SoC rule improves the safety margin, it does so by adding unnecessary charging time relative to the route-adaptive optimum. This illustrates that fixed-SoC rules may be feasible, but they do not provide a systematic way to avoid overcharging across routes with different energy-demand profiles.

The full-charge heuristic is the most conservative alternative, but also the least efficient in terms of mission time. Although it reduces the final EPA probability to 0.102 for Route #1 and 0.156 for Route #2, it increases the charging time to 0.981 h for both routes. Compared with the proposed method, this represents an additional charging time of approximately 0.694 h for Route #1 and 0.664 h for Route #2. Therefore, charging to full capacity reduces risk, but at the expense of a substantial increase in total mission time.

The comparison with simple charging heuristics further showed that the proposed method achieves feasibility without resorting to unnecessary conservative charging. In particular, departing without additional charging led to infeasible missions, while fixed-SoC and full-charge strategies reduced risk at the expense of additional charging time. This highlights the value of adapting the charging target to the route-specific risk profile.

Overall, this baseline comparison complements the route-level optimization results reported in Table 5. While Table 5 identifies the best route under the proposed risk-aware strategy, Table 6 shows why the optimized charging decision is preferable to common heuristic rules. The proposed method achieves feasibility without overcharging, providing the minimum route-adaptive charging action required to satisfy the prescribed EPA risk tolerance.

4.6. Reproducibility and Generalization Limits

To improve the reproducibility of the proposed case study, Table 7 summarizes the main experimental settings, model configurations, and uncertainty-evaluation protocol used in the Costa Rica validation. These details are reported to-

Table 6. Comparison between the proposed risk-aware strategy and simple charging heuristics, considering a pre-charging SoC of $SoC_0 = 55\%$.

Strategy	Route	Departure SoC [%]	Charging time [h]	Total mission time [h]	Final $P(EPA)$	Feasible
Proposed method	Route #1	76.59	0.287	2.286	0.389	Yes
Proposed method	Route #2	79.84	0.317	2.421	0.467	Yes
No additional charging	Route #1	55.00	0.000	1.999	0.915	No
No additional charging	Route #2	55.00	0.000	2.103	0.976	No
Fixed SoC target 85%	Route #1	85.00	0.430	2.429	0.349	Yes
Fixed SoC target 85%	Route #2	85.00	0.430	2.533	0.439	Yes
Full-charge heuristic	Route #1	100.00	0.981	2.980	0.102	Yes
Full-charge heuristic	Route #2	100.00	0.981	3.084	0.156	Yes

gether to make explicit the operational scope of the evaluation, including the telemetry dataset, battery-parameter identification procedure, GRU-LGBM demand-prediction models, charging-optimization settings, and uncertainty metrics.

The GRU and LGBM configurations clarify the predictive stage used before the MDR-UELF risk propagation. The GRU model is responsible for capturing the sequential behavior of vehicle speed along the route, whereas the LGBM model uses route, vehicle-dynamics, and environmental features to estimate segment-level energy and power demand. These predicted demand profiles are then propagated through the battery model to obtain the voltage trajectories and EPA probabilities used by the charging optimization. Therefore, the reported hyperparameters and input features define the data-driven component that connects route information with the final risk-aware decision.

The uncertainty intervals reported in Table 4 were evaluated through the interval coverage metric, defined here as the fraction of measured SoC samples that remained within the predicted MDR-based uncertainty envelope along each route. Therefore, this metric should be interpreted as an empirical consistency indicator between the predicted uncertainty bounds and the battery depletion observed during real operation. In the analyzed case study, the interval coverage reached 1.0000 for Route #1 and 0.9549 for Route #2, indicating that the uncertainty bands were informative for the two evaluated routes.

At the same time, the generalization scope of the present validation must be interpreted with care. The objective of this paper is to demonstrate the feasibility of the proposed risk-aware charging and routing framework under real operating conditions, rather than to claim universal performance across all EV models, battery chemistries, charging stations, or geographical regions. The experimental validation is based on one vehicle and two candidate routes, which provides a controlled and physically interpretable case study, but does not fully characterize fleet-level or multi-region generalization. Extending the evaluation to additional vehicles, routes, charg-

ers, traffic conditions, and battery-aging levels is therefore an important direction for future work.

Nevertheless, the proposed framework is not vehicle-specific in its formulation. The vehicle-dependent elements are mainly the identified battery-model parameters, the charging protocol, and the trained demand-prediction models. Once these components are re-identified or retrained for a new EV-charger pair and operational region, the same MDR-UELF risk evaluation and chance-constrained charging optimization can be applied without modifying the decision-making structure of the method.

In summary, the results show that the proposed framework is not limited to estimating a charging time in isolation. Instead, it connects the charging decision with route-dependent battery depletion, mission-level risk, and total travel duration. Under this perspective, the methodology provides a more operationally meaningful basis for EV route selection, since it allows the decision-maker to choose not only a feasible route, but also a charging-routing strategy that balances time efficiency and probabilistic safety.

5. CONCLUSION

This work presented a risk-aware framework for jointly determining charging time and route selection for electric vehicles under uncertainty. The proposed methodology integrates a realistic EV-charger charging protocol with an MDR-based probabilistic feasibility evaluation, allowing the charging decision to be explicitly linked to route-dependent energy demand, battery depletion, and mission-level risk. Under this formulation, the problem is cast explicitly as a decision-making task in which the user must choose both how much to charge and which route to take while accounting for uncertainty and individual risk tolerance.

The results obtained in the Costa Rica case study show that the framework is able to identify route-specific departure SoC requirements and their associated charging times under a prescribed admissible risk threshold. In particular, the two can-

Table 7. Summary of the main experimental settings, model configurations, and reproducibility details used in the Costa Rica case study.

Category	Item	Setting / Description
Case study	Vehicle and context	2018 Nissan Leaf, real-world operation in Costa Rica
	Telemetry dataset	Sept. 2025–Mar. 2026; ~128,000 raw samples; 107,216 valid samples; 488 trips
	Data split	70% training, 15% validation, 15% test by trip index
Battery model	Model structure	Thevenin-equivalent model with OCV–SoC characterization
	Parameter identification	PSO with 100 particles, 200 iterations, $\omega = 0.72$, $c_1 = c_2 = 1.49$
Demand prediction	GRU speed model	2 layers, 40 hidden units, learning rate 10^{-3} , batch size 128
	GRU input features	Hour, trip distance, node distance, slope, street category
	LGBM energy/power model	Max. depth 120, 3000 estimators, 600 leaves, learning rate 0.1, L1 regularization 0.05
	LGBM input features	Vehicle speed, previous speed, next speed, trip distance, delta distance, slope, humidity, temperature, precipitation
Optimization and uncertainty	Model outputs	Speed trajectory, segment-level energy consumption, and power demand
	Charging optimization	Bisection search over departure SoC; max. 50 bisections; tolerance 1.0×10^{-4} h
	Risk setting	Admissible EPA risk tolerance $\delta = 0.5$; pre-charging SoC $SoC_0 = 55\%$
	Uncertainty propagation	MDR-based stochastic prediction intervals
Evaluation scope	Interval coverage	Fraction of measured SoC samples contained within the predicted interval
	Candidate routes	Two real routes with common origin and destination
	Generalization scope	One vehicle and two routes; extension to additional EVs, chargers, and regions left for future work

didate routes required different departure SoC targets despite connecting the same origin and destination, confirming that the proposed approach is sensitive to route-dependent operating conditions rather than relying on fixed charging heuristics. Moreover, the real-operation validation showed that the predicted SoC trajectories remained consistent with the measured battery depletion trends, while the quantitative error metrics confirmed low prediction error and strong uncertainty coverage in both routes.

From the decision-making perspective, the proposed methodology proved capable of selecting the route that offers the best trade-off between charging requirements, travel time, and probabilistic feasibility. In the analyzed case, the descent route achieved the lowest charging time, the minimum total mission time, and the lowest final EPA probability among the feasible alternatives. These results support the practical value of the framework as an operational tool for EV users facing route and charging decisions in realistic fast-charging scenarios.

Overall, the main contribution of this paper is to move beyond deterministic range heuristics by introducing a unified decision framework in which charging duration and route selection are optimized jointly under uncertainty. In this way, the methodology provides not only a feasible charging recommendation, but also an interpretable notion of route-level risk and its accumulation along the mission.

Future work will focus on extending the framework to scenarios with a larger number of candidate routes and charging stations, explicitly incorporating waiting times and station congestion, and accounting for additional uncertainty sources such as traffic evolution, weather conditions, and battery aging. Another relevant extension is the incorporation of online re-planning capabilities, so that the charging and routing de-

cision can be updated in real time as new operational information becomes available.

ACKNOWLEDGMENT

This work was supported in part by ANID FONDECYT 1250036, Advanced Center for Electrical and Electronic Engineering, ANID Basal Project CIA250006. The work of Jorge E. García Bustos has been supported by ANID-PFCHA/Doctorado Nacional/2022-21221213. The work of Benjamin Brito Schiele has been funded by the IntelliWind Doctoral Network, granted under the Horizon Europe MSCA Doctoral Network programme, Grant agreement 101168725. The authors would also like to thank Mrs. Morelia Soto, Mr. Emanuel Coto, and Mr. Omar Lopez for driving during the experiments used in the case study presented in this work.

REFERENCES

- Acuña-Ureta, D. E., & Orchard, M. E. (2023, 4). Near-instantaneous battery end-of-discharge prognosis via uncertain event likelihood functions. *ISA Transactions*, 135, 199-212. doi: 10.1016/j.isatra.2022.09.040
- Autoridad Reguladora de los Servicios Públicos (ARESEP). (2023). *RE-0020-IE-2023: Fijación de oficio de la tarifa aplicable en los centros de recarga rápida para vehículos eléctricos (T-VE) por tiempo de uso* (Tech. Rep.). Intendencia de Energía, ARESEP. Retrieved from <https://asomove.org/redirect/legislacion/fcle3> (Accessed: 2026-01-22)
- Baptista, M. L., Mishra, M., Henriques, E., & Prendinger, H. (2024). Using explainable artificial intelligence to interpret remaining useful life estimation with gated recurrent unit. In *Annual conference of the phm society*

- (Vol. 16). doi: 10.36001/phmconf.2024.v16i1.4124
- Biresselioglu, M. E., Kaplan, M. D., & Yilmaz, B. K. (2018, 3). Electric mobility in europe: A comprehensive review of motivators and barriers in decision making processes. *Transportation Research Part A: Policy and Practice*, 109, 1-13. doi: 10.1016/j.tra.2018.01.017
- Brown, A., Cappellucci, J., Gaus, M., & Buleje, H. (2024, 11). *Electric vehicle charging infrastructure trends from the alternative fueling station locator: Second quarter 2024* (Tech. Rep.). National Renewable Energy Laboratory (NREL). doi: 10.2172/2478445
- Burgos-Mellado, C., Orchard, M. E., Kazerani, M., Cárdenas, R., & Sáez, D. (2016, 1). Particle-filtering-based estimation of maximum available power state in lithium-ion batteries. *Applied Energy*, 161, 349-363. doi: 10.1016/j.apenergy.2015.09.092
- 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: 10.1016/j.engappai.2024.109925
- Bustos, J. E. G., Schiele, B. B., Masserano, B., Salas-Espiñeira, R., Troncoso-Kurtovic, D., Acuña-Ureta, D. E., ... Perez, A. (2026, 6). Enabling online maximum driving range prognostics in electric vehicles via uncertain event likelihood functions. *Engineering Applications of Artificial Intelligence*, 173, 114449. doi: 10.1016/j.engappai.2026.114449
- Chen, J., Jing, H., Chang, Y., & Liu, Q. (2019). Gated recurrent unit based recurrent neural network for remaining useful life prediction of nonlinear deterioration process. *Reliability Engineering & System Safety*, 185, 372-382. doi: 10.1016/j.ress.2019.01.006
- Clar-Garcia, D., Fabra-Rodriguez, M., Campello-Vicente, H., & Velasco-Sanchez, E. (2025, 10). Optimal dc fast-charging strategies for battery electric vehicles during long-distance trips. *Batteries*, 11, 394. doi: 10.3390/batteries11110394
- Egbue, O., & Long, S. (2012, 9). Barriers to widespread adoption of electric vehicles: An analysis of consumer attitudes and perceptions. *Energy Policy*, 48, 717-729. doi: 10.1016/j.enpol.2012.06.009
- Gnann, T., Funke, S., Jakobsson, N., Plötz, P., Sprei, F., & Bennehag, A. (2018, 7). Fast charging infrastructure for electric vehicles: Today's situation and future needs. *Transportation Research Part D: Transport and Environment*, 62, 314-329. doi: 10.1016/j.trd.2018.03.004
- gridX. (2025). *European ev charging report 2025*. Retrieved 2026-01-22, from www.gridx.ai/resources/european-ev-charging-report-2025 (gridX GmbH, Aachen/Munich, Germany)
- Hanig, L., Ledna, C., Nock, D., Harper, C. D., Yip, A., Wood, E., & Spurlock, C. A. (2025, 1). Finding gaps in the national electric vehicle charging station coverage of the united states. *Nature Communications*, 16, 561. doi: 10.1038/s41467-024-55696-8
- Hoen, F. S., Díez-Gutiérrez, M., Babri, S., Hess, S., & Tørset, T. (2023, 9). Charging electric vehicles on long trips and the willingness to pay to reduce waiting for charging. stated preference survey in norway. *Transportation Research Part A: Policy and Practice*, 175, 103774. doi: 10.1016/j.tra.2023.103774
- Huang, C. H., Soto, M., Ortiz, J., Arguello, A., & Perez, A. (2024, 11). A quick overview of how electric vehicles are rapidly charged. In *2024 ieee pes generation, transmission and distribution latin america conference and industrial exposition (gtla)* (p. 1-5). IEEE. doi: 10.1109/GTDLA61236.2024.10913612
- International Energy Agency (IEA). (2025, may). *Global ev outlook 2025* (Tech. Rep.). Paris: International Energy Agency. Retrieved 2026-01-22, from www.iea.org/reports/global-ev-outlook-2025 (Published 14 May 2025. Licence: CC BY 4.0)
- Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., ... Liu, T.-Y. (2017). Lightgbm: a highly efficient gradient boosting decision tree. In *Proceedings of the 31st international conference on neural information processing systems* (p. 3149-3157). Red Hook, NY, USA: Curran Associates Inc.
- Liu, H., Xiao, Q., Jin, Y., Mu, Y., Meng, J., Zhang, T., ... Teodorescu, R. (2022). Improved lightgbm-based framework for electric vehicle lithium-ion battery remaining useful life prediction using multi health indicators. *Symmetry*, 14(8), 1584. doi: 10.3390/sym14081584
- Liu, R., He, G., Wang, X., Mallapragada, D., Zhao, H., Shao-Horn, Y., & Jiang, B. (2024, 1). A cross-scale framework for evaluating flexibility values of battery and fuel cell electric vehicles. *Nature Communications*, 15, 280. doi: 10.1038/s41467-023-43884-x
- Löbel, F., Borndörfer, R., & Weider, S. (2023). Non-linear charge functions for electric vehicle scheduling with dynamic recharge rates. In *23rd symposium on algorithmic approaches for transportation modelling, optimization, and systems (atmos 2023)* (Vol. 115, pp. 15:1-15:6). Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi: 10.4230/OASIS.ATMOS.2023.15
- Nemirovski, A., & Shapiro, A. (2007, 1). Convex approximations of chance constrained programs. *SIAM Journal on Optimization*, 17, 969-996. doi: 10.1137/050622328
- Noel, L., de Rubens, G. Z., Kester, J., & Sovacool, B. K. (2020, 3). Understanding the socio-technical nexus of nordic electric vehicle (ev) barriers: A qualitative discussion of range, price, charging and knowledge. *En-*

- ergy Policy*, 138, 111292. doi: 10.1016/j.enpol.2020.111292
- Perez, A., Jaramillo, F., Baeza, C., Valderrama, M., Quintero, V., & Orchard, M. (2021, 11). A particle-swarm-optimization-based approach for the state-of-charge estimation of an electric vehicle when driven under real conditions. *Annual Conference of the PHM Society*, 13. doi: 10.36001/phmconf.2021.v13i1.3013
- Sohil, F., Sohali, M. U., & Shabbir, J. (2022, 1). An introduction to statistical learning with applications in r. *Statistical Theory and Related Fields*, 6, 87-87. doi: 10.1080/24754269.2021.1980261
- Tomaszewska, A., Chu, Z., Feng, X., O’Kane, S., Liu, X., Chen, J., . . . Wu, B. (2019, 8). Lithium-ion battery fast charging: A review. *eTransportation*, 1, 100011. doi: 10.1016/j.etrans.2019.100011
- Unterluggauer, T., Rich, J., Andersen, P. B., & Hashemi, S. (2022, 5). Electric vehicle charging infrastructure planning for integrated transportation and power distribution networks: A review. *eTransportation*, 12, 100163. doi: 10.1016/j.etrans.2022.100163
- Zhan, W., Liao, Y., Deng, J., Wang, Z., & Yeh, S. (2025, 3). Large-scale empirical study of electric vehicle usage patterns and charging infrastructure needs. *npj Sustainable Mobility and Transport*, 2, 9. doi: 10.1038/s44333-024-00023-3
- Zhang, H., & Chow, M.-Y. (2010, 11). On-line phev battery hysteresis effect dynamics modeling. In *Iecon 2010 - 36th annual conference on ieee industrial electronics society* (p. 1844-1849). IEEE. doi: 10.1109/IECON.2010.5675395