



Combined Stochastic and Robust Optimization for Electric Autonomous Mobility-on-Demand with Nested Benders Decomposition

Downloaded from: <https://research.chalmers.se>, 2026-08-28 11:39 UTC

Citation for the original published paper (version of record):

Tingstad Jacobsen, S., Kulcsár, B., Lindman, A. (2026). Combined Stochastic and Robust Optimization for Electric Autonomous Mobility-on-Demand with Nested Benders Decomposition. *Transportation Research, Part C: Emerging Technologies*, 192. <http://dx.doi.org/10.1016/j.trc.2026.105903>

N.B. When citing this work, cite the original published paper.



Combined stochastic and robust optimization for electric autonomous mobility-on-demand with nested Benders decomposition

Sten Elling Tingstad Jacobsen ^{a,b,1,*}, Balázs Kulcsár^a, Anders Lindman ^b

^a Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden

^b Volvo Cars, Gothenburg, Sweden

ARTICLE INFO

Keywords:

Nested benders decomposition
Electric autonomous mobility-on-demand
Bayesian neural networks
Scenario tree
Scenario reduction
Charging scheduling
Connected and autonomous vehicles

ABSTRACT

The electrification and automation of mobility are reshaping how cities operate on-demand transport systems. Managing Electric Autonomous Mobility-on-Demand (EAMoD) fleets effectively requires coordinating dispatch, rebalancing, and charging decisions under multiple uncertainties, including travel demand, travel time, energy consumption, and charger availability. We address this challenge with a combined stochastic and robust model predictive control (MPC) framework. The framework integrates spatio-temporal Bayesian neural network forecasts with a multi-stage stochastic optimization model, formulated as a large-scale mixed-integer linear program. To ensure real-time applicability, we develop a tailored Nested Benders Decomposition that exploits the scenario tree structure and enables efficient parallelized solution. Stochastic optimization is employed to anticipate demand and infrastructure variability, while robust constraints on energy consumption and travel times safeguard feasibility under worst-case realizations. We evaluate the framework using high-fidelity simulations of San Francisco and Chicago. Compared with deterministic, reactive, and robust baselines, the combined stochastic and robust approach reduces median passenger waiting times by up to 36% and 95th-percentile delays by nearly 20%, while also lowering rebalancing distance by 27% and electricity costs by close to 35%. We also conduct a sensitivity analysis of battery size and vehicle efficiency, finding that energy-efficient vehicles maintain stable performance even with small batteries, whereas less efficient vehicles require larger batteries and greater infrastructure support. Our results emphasize the importance of jointly optimizing predictive control, vehicle capabilities, and infrastructure planning to enable scalable, cost-efficient EAMoD operations.

1. Introduction

As cities worldwide face rapid urban growth, environmental challenges, and evolving societal demands, transportation systems play a critical role in achieving sustainable urban development. Expanding urban populations worsen issues such as traffic congestion, pollution, aging infrastructure, and socioeconomic inequality, threatening the sustainability and livability of urban areas (Chee and Neo, 2018; United Nations Development Programme, 2017). In response, the concept of smart cities has emerged, leveraging

* Corresponding author.

E-mail addresses: elling@chalmers.se (S.E. Tingstad Jacobsen), kulcsar@chalmers.se (B. Kulcsár), anders.lindman@volvocars.com (A. Lindman).

¹ Principal corresponding author.

advanced technologies such as Intelligent Transportation Systems (ITS). ITS integrate real-time data, connected infrastructure, and autonomous vehicles to enhance urban management, optimize mobility, and promote sustainable transportation solutions (Nature Editorial; United Nations Development Programme (UNDP), 2020; Batty et al., 2012). A promising development within this domain is Electric Autonomous Mobility-on-Demand (EAMoD) systems. These systems consist of centrally coordinated fleets of electric autonomous vehicles that provide on-demand transport services (Spieser et al., 2014). EAMoD systems have the potential to reduce urban emissions, improve operational efficiency, and accelerate the adoption of electric vehicles (EVs). Unlike fixed-route public transport, EAMoD offers dynamic routing and centralized coordination, enabling adaptation to real-time demand (Zardini et al., 2022).

Optimizing such systems requires solving several interdependent problems: dispatching vehicles to satisfy spatially and temporally varying demand, rebalancing empty vehicles to anticipate future imbalances, and scheduling charging and vehicle-to-grid (V2G) operations under infrastructure and energy constraints. These decisions must be made jointly and under uncertainty, as travel demand, charger availability, travel time, and energy consumption are all stochastic and temporally correlated. Existing approaches typically address only a subset of these challenges. Many studies focus on vehicle allocation without considering the complexities introduced by electric vehicles, such as limited battery range and charging infrastructure constraints (Zardini et al., 2022). Those that incorporate charging often treat it separately from dispatch and rebalancing, or rely on deterministic forecasts that fail to account for the full range of future uncertainty (He et al., 2023; Gao and Li, 2024). On the scenario modeling side, constructing reduced scenario trees that faithfully represent multivariate, temporally correlated uncertainties remains computationally demanding, and most existing reduction techniques are limited to single sources of uncertainty or do not preserve the filtration structure of the stochastic process (Horejšová et al., 2020). Finally, solving the resulting multi-stage stochastic optimization problems at a scale suitable for real-time operations is itself a major challenge, as popular decomposition methods such as SDDP assume stagewise independence, an assumption violated in EAMoD applications where uncertainties exhibit strong temporal dependencies (Zou et al., 2019).

In this paper, we propose a combined stochastic and robust model predictive control (MPC) framework for EAMoD fleet optimization that addresses these challenges in an integrated manner. The framework uses Bayesian Neural Networks (BNNs) (Saad et al., 2024) to produce calibrated probabilistic forecasts of travel demand and charger availability, capturing both spatial and temporal correlations. These forecasts are used to generate a multi-stage scenario tree, which is then reduced using a clustering-based technique that efficiently handles multiple sources of uncertainty. The resulting large-scale Mixed-Integer Linear Program (MILP) is solved using a tailored Nested Benders Decomposition (NBD) algorithm that exploits the scenario tree structure and enables parallelized computation. Robust chance constraints on energy consumption and travel time safeguard feasibility under worst-case realizations. The framework is evaluated through high-fidelity simulations in San Francisco and Chicago, using real-world taxi data, and is complemented by a comprehensive sensitivity analysis of vehicle efficiency, battery capacity, fleet size, and charging infrastructure. An overview of the framework is presented in Fig. 1.

The main contributions of this paper are:

- A combined stochastic and robust MPC framework for EAMoD that jointly optimizes dispatch, rebalancing, charging, and V2G operations under four sources of uncertainty: travel demand, charger availability, energy consumption, and travel time.
- A BNN-based probabilistic forecasting pipeline integrated with a clustering-based scenario tree generation and reduction method that preserves temporal correlations across multiple uncertainty sources while maintaining computational tractability.
- A tailored Nested Benders Decomposition (NBD) algorithm that exploits the scenario tree structure to efficiently solve the resulting large-scale multi-stage stochastic MILP, enabling real-time applicability with parallelized subproblem solution.
- A comprehensive evaluation on two large-scale urban case studies (San Francisco and Chicago) demonstrating significant improvements over deterministic, reactive, and robust baselines, including up to 36% reduction in median wait times, 20% reduction in 95th-percentile delays, 27% less rebalancing distance, and close to 35% savings in electricity costs.
- A sensitivity analysis investigating the interplay between vehicle energy efficiency, battery capacity, fleet size, and charger availability, providing practical insights for fleet design and infrastructure planning in EAMoD systems.

Importantly, the proposed framework differs from existing multi-stage approaches in two key aspects. First, the uncertainty process is temporally dependent and multivariate, violating the stagewise independence assumption required by SDDP-type methods. Second, the integration of robust chance constraints within a multi-stage SMPC architecture enables feasibility guarantees while retaining anticipatory decision-making. These features necessitate a tailored Nested Benders Decomposition capable of handling large scenario trees with temporal coupling.

The remainder of this work is organized as follows. Section 2 reviews the related literature on EAMoD optimization, scenario tree methods, and decomposition algorithms. Section 3 presents the system model and the complete multi-stage robust and stochastic MPC formulation. Section 4 describes the BNN-based scenario generation and clustering-based reduction procedure. Section 5 presents the NBD algorithm. Section 6 reports simulation results for San Francisco and Chicago, including a sensitivity analysis, and concludes with a discussion of findings and future research directions.

2. Related work

2.1. AMoD demand modeling and uncertainty quantification

Optimizing EAMoD systems introduces several challenges, particularly in addressing demand-supply imbalances resulting from spatial and temporal variations in travel demand. Numerous studies have proposed various vehicle allocation control methods to

mitigate these imbalances. For instance, Zhang et al. (2016) introduced a Model Predictive Control (MPC) approach that demonstrated promising results; however, its scalability was limited due to the need to model the entire road network. To enhance scalability, many studies have adopted discretization strategies, dividing the city into larger areas, often referred to as stations. An example of this approach is presented in Iglesias et al. (2019), where a queue-theoretical model was employed. An alternative and effective model-free approach is the "+1 method," which directs vehicles to the most recent service request location, as demonstrated by Ruch et al. (2020). For MPC methods, incorporating uncertainties into the control algorithm has proven important. A stochastic MPC framework combined with a Long Short-Term Memory (LSTM) neural network for travel demand prediction, leveraging Sample Average Approximation (SAA) to model uncertainty, was proposed by Tsao et al. (2018). Similarly, Jacobsen et al. (2023) introduced a chance-constrained approach using Gaussian Process Regression (GPR) for demand prediction, while Miao et al. (2021) proposed a data-driven distributionally robust optimization method. Despite these advancements, many existing strategies often overlook the complexities associated with electric vehicles (EVs), including limited battery capacities and charging requirements, as highlighted in the review by Zardini et al. (2022).

To address this gap, several studies have explored charge scheduling strategies in EAMoD. Charging strategies such as plug-in versus battery swapping have been investigated to optimize fleet availability and operational efficiency in EAMoD systems (Gao and Li, 2024). The placement of charging stations using stochastic programming and considering different uncertainties was studied in Faridimehr et al. (2018). The charging dynamics based on data from an EV taxi fleet were examined in Lei et al. (2022). Fleet sizing and charging station placement considering endogenous congestion were analyzed in Yang et al. (2023). While these studies address important aspects of EAMoD fleets, the challenging task of simultaneously considering charge scheduling and vehicle allocation is not explored. Some of these aspects are considered in He et al. (2023), where a distributionally robust approach was proposed to account for uncertainties in demand and supply. An ARIMA model was used for travel demand prediction, and the uncertainty set was constructed using bootstrapping. Different reinforcement learning algorithms have also been proposed. A queue-theoretical deep reinforcement learning method incorporating routing, battery charging, and dynamic pricing was introduced by Turan et al. (2020). Various reinforcement learning strategies are compared in the survey by Qin et al. (2022). However, a comprehensive approach that simultaneously addresses both charging scheduling and vehicle allocation while accounting for uncertainties in travel demand, charging availability, energy consumption and travel time remains unexplored.

2.2. Scenario generation, reduction, and quality metrics

A common approach to handling uncertainties in control and optimization is to construct scenario trees. A scenario tree provides a structured representation of the evolution of uncertainty over time in a multi-stage decision process. Each node in the tree corresponds to a possible future state, while the branches represent different ways in which uncertainty can unfold, along with their associated probabilities. It has been demonstrated that increasing the number of scenarios in the tree improves the accuracy of the uncertainty representation. However, it should be noted that an increase in the number of scenarios also results in an increase in the computational complexity of the optimization problem Beltrán et al. (2021). In order to address this issue, a range of scenario reduction techniques have been developed with the objective of constructing a reduced scenario tree that closely approximates the original one. These methods typically rely on a metric that quantifies the difference between the original and reduced trees, which can be used both to guide the reduction procedure and to evaluate its quality. A prevalent methodology in this field is moment matching, whereby the reduced scenario tree is constructed in order to preserve the first few moments (e.g., mean, variance, skewness, and kurtosis) of the original stochastic process (Høyland and Wallace, 2001). This technique has been employed in several studies to generate scenario trees by matching the first four moments (Ansariipoor et al., 2016; Bao et al., 2023). However, it has been shown that two probability distributions with matching moments can still be significantly different (Hochreiter and Pflug, 2007). An alternative metric is the Wasserstein distance, which quantifies the similarity between two probability distributions and has been used for scenario reduction (Dupačová et al., 2003). However, the Wasserstein distance does not consider the filtration structure of the scenario tree, meaning it does not account for the temporal correlations between different stages. To address this, the Nested distance, also referred to as process distance, was introduced in Pflug (2010). The Nested distance explicitly accounts for the tree structure, preserving the temporal dependencies of the stochastic process. Although computing the Nested distance between two scenario trees is relatively straightforward, finding an optimally reduced scenario tree that minimizes this distance is computationally challenging. One approach to scenario reduction based on the Nested distance involves multistage k-means clustering, as proposed in Kovacevic and Pichler (2015). However, this method still requires solving multiple large-scale linear programs (LPs), making it computationally demanding. A comprehensive evaluation of different scenario reduction techniques using the Nested distance is provided in Horejšová et al. (2020), which highlights that methods minimizing the Nested distance tend to be computationally expensive. Another challenge is reducing scenario trees that incorporate multiple sources of uncertainty.

2.3. Multistage stochastic optimization and decomposition methods

Despite advances in scenario reduction, solving multi-stage stochastic programs (SPs) remains computationally challenging, particularly as the number of scenarios and the planning horizon increase. To address this, a wide range of decomposition techniques have been developed. The classical Benders decomposition method (Benders, 1962) laid the foundation for solving mixed integer linear programs (MILPs) with complicating variables, and was extended into the L-shaped method (Van Slyke and Wets, 1969) for two-stage SPs. Subsequent work introduced regularized variants (Ruszczyński, 1986) and stochastic decomposition methods (Higle and Sen, 1991; Sen and Higle, 1999) inspired by stochastic quasi-gradient approaches (Ermoliev, 1983) and stochastic approximation

(Robbins and Monro, 1951). For multistage problems, nested Benders decomposition (NBD) (Birge, 1985) and multistage stochastic decomposition (Sen and Zhou, 2014) were developed to extend these ideas. A major advancement came with sampling-based methods such as stochastic dual dynamic programming (SDDP) (Pereira and Pinto, 1991), and its subsequent enhancements (Infanger and Morton, 1996; Shapiro et al., 2011), as well as the stochastic dual dynamic integer programming (SDDiP) approach (Zou et al., 2019). Alternative strategies include abridged nested Benders (Donohue et al., 2006) and cutting-plane or partial sampling methods (Chen and Powell, 1999). SDDP and SDDiP are especially effective for convex or mixed-integer multistage problems but typically rely on the assumption of stagewise independence of the underlying uncertainty process, which facilitates cut sharing across nodes at the same stage and enables tractable backward recursion (Shapiro et al., 2011; van Ackooij and Warin, 2020). While extensions exist that embed dependence via additional state variables, the stagewise independence assumption remains central to the classical framework. This assumption can be restrictive in real-world applications where uncertainties exhibit strong temporal correlations. In EAMoD systems, travel demand, energy consumption, and charger availability evolve with significant temporal dependence, violating the structural assumptions underlying standard SDDP-type methods and motivating alternative decomposition strategies that explicitly operate on temporally dependent scenario trees.

In transportation research, variants of Benders decomposition have been successfully applied to a variety of different problems. A stabilized Benders algorithm has been used to solve a two-stage stochastic program for locating recharging stations in one-way electric car-sharing systems under demand uncertainty, demonstrating scalability to hundreds of large-scale demand scenarios (Çalik and Fortz, 2019). An exact Benders decomposition embedded in a branch-and-cut framework has been developed to strategically locate mini-depots for last-mile crowd-shipping, achieving significant computational speedups and cost savings (Nieto-Isaza et al., 2022). A hybrid approach combining Benders decomposition with NSGA-II has been proposed for the joint optimization of EV charging station placement and pricing, enabling efficient infrastructure deployment while balancing conflicting objectives such as consumer pricing and operator profitability (Ameer et al., 2025). Despite these advances, existing decomposition methods either assume stagewise independence (e.g., SDDP), rely on convexity, or are not designed for real-time MPC implementations with integer variables and temporally correlated uncertainties. The EAMoD setting considered here requires solving large-scale multi-stage MILPs with coupled uncertainties and integrality constraints, motivating the development of a tailored Nested Benders approach.

The fleet composition, including the fleet size and individual vehicle attributes such as battery size and energy efficiency, affects how the fleet is optimized, the service quality provided, and the operational cost. For example, a less energy-efficient vehicle would need to charge more often compared to an efficient vehicle. Alternatively, the less energy-efficient vehicle could have a larger battery to compensate for its inefficiency, but this would increase costs and have a higher environmental impact. Few studies have incorporated fleet composition into the evaluation of EAMoD fleet performance. The effects of battery size and fleet size have been explored in Paparella et al. (2024). However, fast charging and varying energy efficiency were not considered.

3. EAMoD vehicle movement modeling

The model presented in this section is designed for centralized control of electric mobility-on-demand systems, focusing on optimizing vehicle movement and bi-directional charging scheduling to minimize costs while maximizing service completion levels. These high-level decisions serve as references for low-level controllers, which make real-time adjustments. Central to our model is its predictive nature across time, space, and energy dimensions. By exploiting its predictive capabilities, we aim to make optimal decisions in the current time step based on expected future states evolving in time and space. The model is tailored for receding horizon optimization, ensuring adaptability to changing conditions.

Our model is comprised of two integral components. First, we build upon the network model previously established in Jacobsen et al. (2023) by incorporating electric vehicles (EVs) into the framework. This enhancement requires adding several key features specific to EVs. First, charging stations to accommodate the recharging of vehicles, second the limited range of EVs necessitating strategic planning around vehicle deployment, and third vehicle-to-grid (V2G) capabilities, which allow for the interaction between the vehicle's battery and the power grid for energy exchange. Secondly, we introduce uncertainties into the model, encompassing aspects such as demand variability, charging station availability, travel time, and energy consumption rates.

3.1. Model notation

The Electric Autonomous Mobility on Demand (EAMoD) framework represents the city as comprising n distinct stations or areas, enumerated as $\mathcal{N} = [1, \dots, n]$. Stations are identified by applying a K-means clustering algorithm to travel data, (Jin and Han, 2017). The chargers, positioned at each station's centroid, are publicly accessible and the charger occupancy is assumed to follow a non-stationary Poisson process. This simplification avoids the complexities of optimizing charger locations, for which existing literature provides insights (Kchaou-Boujelben, 2021; Yang et al., 2017). Travel is characterized by origin-destination pairs, with the origin denoted by index $i \in \mathcal{N}$ and the destination by index $j \in \mathcal{N}$. The state of charge (SoC) of vehicles is categorized into a set $\mathcal{B} = [0, 0.01, \dots, 0.99, 1]$, defining each SoC level by the index b . Time is discretized into $\mathcal{T} = [1, 2, \dots, T]$, where t is the time index and T the prediction horizon. A summary of different variables used in the model is provided in Table 1.

3.2. Energy, space and time network model

To accurately represent the dynamic environment in which vehicles operate, we introduce an Energy, Space, and Time Network Model. This model forecasts the number of vehicles at the different stations in different time intervals, and their SoC levels, providing

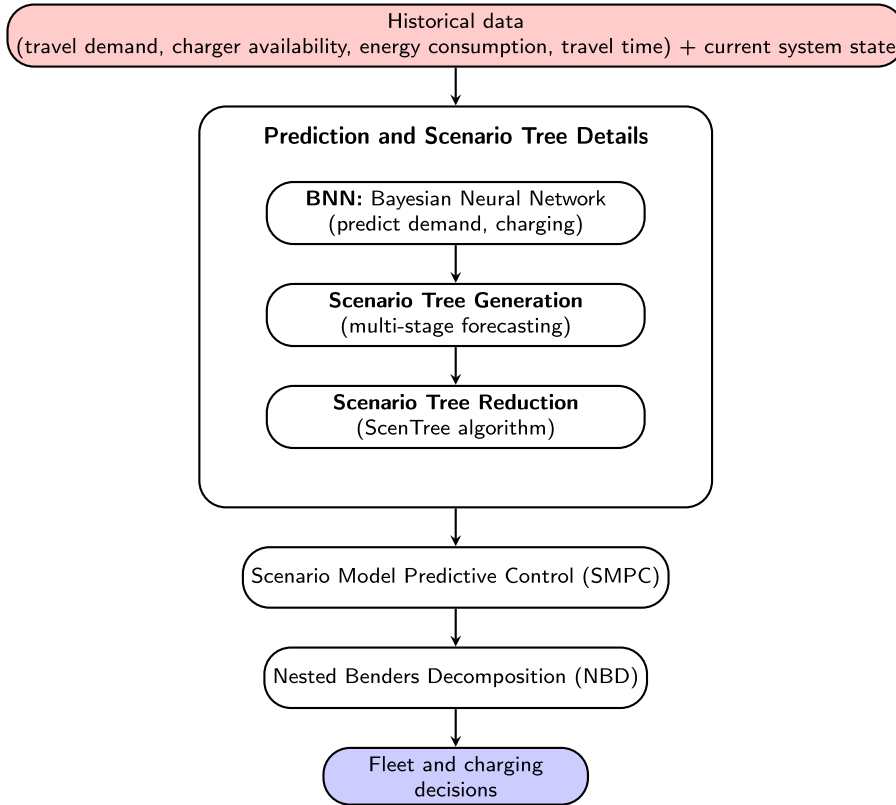


Fig. 1. Overview of the scenario-based model predictive control (SMPC) framework using a Bayesian Neural Network (BNN), scenario tree generation and reduction, and Nested Benders Decomposition (NBD). Historical data is first used to train the BNN, which provides predictions of future demand and charging availability. These predictions are used to generate a multi-stage scenario tree, which is then reduced using the ScenTree algorithm to obtain a computationally tractable representation. The reduced scenario tree is used within the SMPC, which is solved using NBD to derive optimal fleet rebalancing and charging decisions.

a comprehensive understanding of fleet operations (see Fig. 2). Rather than tracking individual vehicle trajectories, the model aggregates vehicles within each station based on identical states. This aggregation enhances computational tractability, enabling efficient analysis of fleet distribution and status across the network while maintaining manageable complexity. The stations are assumed to be fully interconnected, allowing unrestricted travel between any two locations. Consequently, the road network is modeled as a complete graph, ensuring total connectivity among all stations. In this framework, the travel time and energy consumption between any two stations i and j at time t are denoted as Δt_{ijt} and Δb_{ijt} , respectively. These values are time-dependent, as they may vary throughout the day.

Uncertainty at stage t is represented by the random vector ξ_t , with history $\xi_{[t]} = (\xi_1, \dots, \xi_t)$. In the scenario tree approximation, each scenario is indexed by $s \in S$, and the corresponding realization of uncertainty at stage t is denoted ξ_t^s . Decision variables are therefore indexed by s , while ξ_t is reserved exclusively for the underlying random variable. This distinction ensures a clear separation between the stochastic process and its sampled realizations in the Sample Average Approximation (SAA) formulation.

3.2.1. Flow dynamics

There are three types of vehicle movements in the network model. First, passenger transport is represented by x_{ijbt}^p , denoting vehicles traveling from origin i to destination j with SoC b at time t . Second, vehicle relocation is modeled using x_{iibt}^r , which represents vehicles moving between stations without passengers, a crucial mechanism for balancing vehicle distribution across the network. Third, vehicles engaged in charging or V2G operations are represented by x_{iibt}^c and x_{iibt}^d , respectively. The index ii is used for these operations because charging and V2G are modeled as trips within the same station.

The flow conservation constraint tracks vehicle movements across the network over time. At each time step, vehicles may become available at a station from two sources: (i) vehicles that were stationary at the start of the optimization horizon ($t = 0$), representing the initial fleet distribution, and (ii) vehicles that were dispatched before the current horizon began and arrive during the planning period ($t > 0$). We denote by ϕ_{ibt} the number of vehicles becoming available at station i with SoC b at time t . This is an exogenous parameter determined by the current system state: at $t = 0$, it reflects the initial fleet distribution; at $t > 0$, it captures predetermined arrivals from trips initiated in previous receding horizon iterations. These arrivals cannot be altered by the current optimization.

Table 1

Summary of all indices, state variables, decision variables, slack variables, and parameters used in the model.

Category	Symbol	Description
Indices		
	$i \in \mathcal{N}$	Origin of travel
	$j \in \mathcal{N}$	Destination of travel
	$b \in \mathcal{B}$	SoC
	$t \in \mathcal{T}$	Time
	$s \in \mathcal{S}$	Scenario sample
State Variables		
	$\lambda_{ij} \in \mathbb{N}$	Number of passengers traveling from i to j
	$\phi_{ibt} \in \mathbb{N}$	Number of vehicles becoming available at station i with SoC b at time t (exogenous)
	$k_{it}^c \in \mathbb{N}$	Number of available chargers at station i at time t
Decision Variables		
	$x_{ijbt}^c \in \mathbb{N}$	Vehicles charged at station i with SoC b
	$x_{ijbt}^d \in \mathbb{N}$	Vehicles used for V2G at station i with SoC b
	$x_{ijbt}^p \in \mathbb{N}$	Vehicles transporting passengers from i to j with SoC b at time t
	$x_{ijbt}^r \in \mathbb{N}$	Vehicles rebalancing from i to j with SoC b at time t
Slack Variables		
	$s_{ijt} \in \mathbb{N}_{\geq 0}$	Customers not served at station i wishing to travel to j at time t (non-negative)
Parameters		
	Δb^c	SoC charge in one time interval
	Δd	SoC discharge to the grid in one time interval
	e^b	Battery capacity [kWh]
	e_{ijt}	Energy consumption for travel from i to j at time t
	P_b^c	Charging power in kWh for different SoC levels b
	P^d	V2G transfer power in kWh
	Δt_{ijt}	Travel time between station i and j at time t
	Δb_{ijt}	SoC consumption between station i and j at time t
	$b_{ij}^{\min}$	Minimum required SoC to drive a customer from station i and j at time t
	s	Index for different scenarios
	$\mathcal{S}_t$	Set of scenarios at stage t
	$b_{\max}$	Max SoC threshold that is penalized in the terminal cost
	$d_{\max}$	Maximum penalty in terminal cost
	e_{safety}	Safety buffer for energy consumption [kWh]
	$\hat{\tau}_{ijt}$	Conservative travel time estimate for route (i, j) at time t

To ensure conservation of vehicles in the system, we enforce a flow conservation constraint at each station:

$$x_{iibt}^c + x_{iibt}^d + \sum_{j=1}^{\mathcal{N}} [x_{ijbt}^p + x_{ijbt}^r] = \phi_{ibt} + x_{ii(b-\Delta b^c)(t-1)}^c + x_{ii(b+\Delta d)(t-1)}^d + \sum_{j=1}^{\mathcal{N}} [x_{ji(b+\Delta b_{jit})(t-\Delta t_{jit})}^p + x_{ji(b+\Delta b_{jit})(t-\Delta t_{jit})}^r] \quad (1)$$

$$\forall i \in \mathcal{N}, t \in \mathcal{T}, \forall b \in \mathcal{B}.$$

where Δb_{jit} denotes the expected SoC reduction during a trip from station j to station i , and Δb^c and Δd represent energy gained or lost through charging and V2G operations, respectively. The left-hand side represents the outflow of vehicles from station i and the right-hand side the inflow of vehicles to station i . ϕ_{ibt} is the initial condition of the inflow from $t < 0$. At the start of each receding horizon iteration, ϕ_{ibt} is computed from the known positions and scheduled arrival times of all vehicles in the fleet.

The first four terms in Eq. (1) represent vehicles entering or exiting charging or V2G services. The summation terms account for vehicles departing (first two terms) and arriving (last two terms) at each station. This equation ensures system-wide vehicle conservation while incorporating temporal and energy constraints. Specifically, each vehicle movement is subject to a travel time delay Δt_{jit} and an associated SoC reduction Δb_{jit} , given by:

$$\Delta b_{ijt} = \lfloor \frac{e_{ijt}}{e^b} \rfloor,$$

where e^b is the battery capacity and e_{ijt} is the energy consumption for travel between stations i and j at time t . Since energy consumption and travel time are subject to uncertainty (e.g., due to traffic conditions, auxiliary loads, and route variability), the parameters Δb_{ijt} and Δt_{jit} used in the flow conservation constraint (1) must be chosen conservatively to ensure operational feasibility. We determine these values using chance constraints on the underlying random variables. Specifically, the travel time parameter is set to a conservative estimate $\hat{\tau}_{ijt}$ satisfying

$$\Delta t_{jit} = \hat{\tau}_{jit}, \quad \text{where } \mathbb{P}(\tau_{jit} \leq \hat{\tau}_{jit}) \geq 1 - \epsilon^t, \quad (2)$$




Fig. 2. This diagram illustrates a multi-layer network representing an Electric Autonomous Mobility-on-Demand (EAMoD) system. Each layer corresponds to a different battery state of charge (SoC). When a vehicle travels in the network (illustrated by a dashed arrow with a vehicle icon), it consumes energy, leading to a decrease in its state of charge (SoC). In contrast, a dashed arrow with a charging station represents a charging trip, during which the vehicle remains at the station and its SoC increases.

ensuring that the actual travel time does not exceed the value used in the model with probability at least $1 - \epsilon^t$. Similarly, the energy consumption parameter is set to

$$\Delta b_{ijt} = \hat{b}_{ijt}, \quad \text{where } \mathbb{P}(e_{ijt} \leq \hat{b}_{ijt} \cdot e^b + e_{\text{safety}}) \geq 1 - \epsilon^b, \quad (3)$$

where e_{safety} is a safety buffer. These conservative estimates are computed offline at the start of each receding horizon iteration using the empirical distributions of travel time and energy consumption, and are then treated as fixed parameters in the optimization. Smaller values of ϵ^b and ϵ^t correspond to more conservative (robust) policies.

3.2.2. Imbalance dynamics

EAMoD systems are prone to imbalances caused by spatial and temporal variations in travel demand. To mitigate these imbalances through effective vehicle allocation, it is essential to model the imbalance at each station. We define the imbalance at a station as the difference between the number of passengers awaiting service and the number of available vehicles. Ideally, system operations should maintain this imbalance at zero. However, due to demand fluctuations and operational constraints, achieving a perfect balance is often infeasible. Therefore, we introduce the following constraint, which allows for a positive imbalance:

$$\lambda_{ijt} - \sum_{b=b_{ij}^{\min}}^1 x_{ijbt}^p + s_{ij(t-1)} = s_{ijt}, \quad \forall i, j \in \mathcal{N}, \forall t \in \mathcal{T}, \quad (4)$$

where λ_{ijt} denotes the travel demand from station i to station j at time t , and x_{ijbt}^p represents the number of vehicles with a SoC between $b_{ij}^{\min}$ and 1. The threshold $b_{ij}^{\min}$ is the minimum SoC required for a vehicle to complete a trip from i to j with high probability, derived from the energy consumption chance constraint (3):

$$b_{ij}^{\min} = \min \left\{ b \in \mathcal{B} : \mathbb{P}(b \cdot e^b - e_{ijt} \geq e_{\text{safety}}) \geq 1 - \epsilon^b \right\}. \quad (5)$$

This ensures that only vehicles with sufficient charge to complete the trip (including a safety buffer e_{safety}) with probability at least $1 - e^b$ are eligible for dispatch. By restricting the summation in (4) to SoC levels $b \geq b_{ij}^{\min}$, the constraint directly links vehicle dispatch decisions to the robust energy requirements. The variable s_{ijt} is a slack variable that captures unmet demand, thereby allowing some flexibility in maintaining balance.

3.2.3. Charging dynamics

Charging dynamics within the network describe how the SoC of vehicles at each charging node evolves. The increase in SoC, denoted as Δb , during each time interval is calculated by the equation:

$$\Delta b^c = \left\lceil \frac{P^c(b)\Delta t}{e^b} \right\rceil.$$

In this equation, $P^c(b)$ is the SoC-dependent charging power and Δt is the discrete time-step length. The operator $\lceil \cdot \rceil$ rounds Δb^c up to the nearest integer, ensuring discrete SoC bins. Charging power typically exhibits non-linear behavior with SoC rates slow as the battery approaches full charge. We adopt the charging-power curve $P^c(b)$ from Wassiliadis et al. (2021), which remains approximately constant up to an SoC of 45% and then decreases linearly.

Within each discrete time interval, vehicles at a charging station may begin a charging session, continue a session from a previous interval, or complete their charging. We assume a one-interval, Δt , charging duration per decision step. This simplification fits with the receding horizon control strategy used in the model, where decisions are regularly updated. By limiting charging sessions to a single interval, the model remains effective and responsive to changes in charging demand and infrastructure availability, without the need to manage the complexities of prolonged charging durations.

The energy depletion of the battery when the vehicle is V2G is given by:

$$\Delta d = \left\lceil \frac{P^d \Delta t}{e^b} \right\rceil.$$

where P^d is the power delivered to the grid, which is considered to be constant. Additionally, the availability of chargers at each station is limited, which necessitates a constraint on the total number of vehicles that can charge or await charging at any given time:

$$\sum_b (x_{ibt}^c + x_{ibt}^d) \leq k_{it}^c, \quad \forall i \in \mathcal{N}, \forall t \in \mathcal{T}, \quad (6)$$

where k_{it}^c represents the total number of charging stations available at each site. This constraint ensures that the model accurately reflects the actual limitations of the charging infrastructure, preventing over-allocation of charging resources.

3.3. Objective function

We formulate an economic objective function in which each term corresponds to a specific monetary cost or revenue. The goal is to minimize the expected total cost under uncertainty in travel demand, charging-station occupancy, energy consumption, and travel time. The true, unknown, and time-varying probability distributions of these random variables are denoted by D_λ^* , D_k^* , D_b^* , and D_τ^* , respectively. The uncertainties in energy consumption D_b^* and travel time D_τ^* are handled via chance constraints. The distributions D_λ^* and D_k^* enter directly into the expected cost. At each stage $t \in \{1, \dots, T\}$, the exogenous uncertainty is described by the joint random vector

$$\xi_t \triangleq (\{\lambda_{ijt}\}_{i,j=1}^N, \{k_{it}^c\}_{i=1}^N), \quad (7)$$

where $\lambda_{ijt} \geq 0$ is the travel demand from zone i to zone j at stage t , and $k_{it}^c \geq 0$ is the charging-station occupancy at zone i at stage t . We denote the history of observations up to stage t by $\xi_{[t]} \triangleq (\xi_1, \dots, \xi_t)$. The joint conditional distribution of ξ_t given the history $\xi_{[t-1]}$ is denoted $D_t(\cdot \mid \xi_{[t-1]})$.

The multi-stage SMPC objective is:

$$J = \underbrace{\sum_{i=1}^N \sum_{j=1}^N \sum_{b=0}^1 Q(x_{ijb0}^r, x_{ijb0}^p, x_{iib0}^c, x_{iib0}^d)}_{\text{stage } t=0} + \mathbb{E}_{\xi_1 \sim D_1} \left[\sum_{i,j,b} Q_1 + \mathbb{E}_{\xi_2 \sim D_2(\cdot \mid \xi_{[1]})} \left[\sum_{i,j,b} Q_2 + \dots \right. \right. \\ \left. \left. \dots + \mathbb{E}_{\xi_T \sim D_T(\cdot \mid \xi_{[T-1]})} \left[\sum_{i,j,b} Q_T + \Phi(x_{ijbT}^r, x_{ijbT}^p, x_{iibT}^c, x_{iibT}^d) \right] \dots \right] \right], \quad (8)$$

where we use the shorthand $Q_t \triangleq Q(x_{ijbt}^r, x_{ijbt}^p, x_{iibt}^c, x_{iibt}^d)$ for brevity.

The stage cost Q_t is defined as

$$Q(x_{ijbt}^r, x_{ijbt}^p, x_{iibt}^c, x_{iibt}^d) = c_{ijb}^r x_{ijbt}^r - c_{ijb}^p x_{ijbt}^p + c_{ibt}^c x_{iibt}^c - c_{ibt}^d x_{iibt}^d, \quad (9)$$

where the cost coefficients $c_{ijb}^r, c_{ijb}^p, c_{ibt}^c, c_{ibt}^d$ are defined in Table 2.

Table 2
Cost and profit components used in the optimization model.

Symbol	Description
c_{ijb}^r	The cost of rebalancing a vehicle from station i to station j with SoC b at time t , representing the expected energy consumption cost.
c_{ibt}^c	The cost of fast charging a vehicle at station i at time t , where the electricity price may vary depending on the time of day.
c_{ijb}^p	The profit from transporting customers from station i to station j with SoC b at time t . This profit is calculated based on the distance, with a constant cost per distance throughout the day.
c_{ibt}^d	The profit from V2G services provided by a vehicle at station i at time t . The profitability is influenced by the varying electricity prices throughout the day.

The terminal cost Φ penalizes vehicles whose state of charge (SoC) is below the maximum level $b_{\max}$ at the end of the horizon:

$$\Phi = \min(q_{\max}, \max((b_{\max} - b) a, 0)) \left(\sum_{i,j,b} x_{ijbT}^r + x_{ijbT}^p + \sum_{i,b} x_{iibT}^c + x_{iibT}^d \right). \quad (10)$$

This term accounts for continued operation of the EAMoD system beyond the optimization horizon. The constant $q_{\max}$ bounds the penalty, and the parameter a converts the battery shortfall $(b_{\max} - b)$ into a monetary cost per unit of missing energy, proportional to anticipated future service requirements and electricity prices.

The nested expectation structure in (8) reflects the sequential nature of the problem: the cost at stage t is incurred after observing $\xi_{[t]}$ but before $\xi_{[t+1]}$ is revealed, so each inner expectation conditions on the history available at that stage. This is the standard formulation of multi-stage stochastic programming (Shapiro et al., 2011).

A key challenge in EAMoD optimization is the mismatch between short-term trip dynamics (minutes) and longer-term charging economics (hours to a full day). We address this by leveraging day-ahead electricity prices, which are known in advance, to dynamically adjust the terminal cost parameter a at each receding horizon iteration. Let p_t denote the electricity price at time t and let $\mathcal{T}^+ = \{T+1, \dots, T+H\}$ denote the lookahead window beyond the prediction horizon, where H is the lookahead length. The adaptive terminal cost parameter is computed as:

$$a_t = a_{\text{base}} \cdot \frac{\bar{p}^+}{\max(p_t, p_{\min})}, \quad \text{where } \bar{p}^+ = \frac{1}{H} \sum_{\tau \in \mathcal{T}^+} p_\tau. \quad (11)$$

Here, a_{base} is the baseline terminal cost parameter, p_t is the current electricity price, $\bar{p}^+$ is the average electricity price over the lookahead window, and $p_{\min} > 0$ is a minimum price threshold that prevents unbounded terminal cost scaling during periods of very low or zero electricity prices (e.g., due to excess solar generation). When future prices exceed current prices ($\bar{p}^+ > p_t$), the ratio increases a_t , incentivizing immediate charging. Conversely, when future prices are lower ($\bar{p}^+ < p_t$), the reduced a_t encourages deferring charging to take advantage of cheaper electricity. This adaptive terminal cost effectively encodes day-ahead price information into the short-horizon optimization, enabling the SMPC to coordinate charging schedules over daily time scales without explicitly extending the prediction horizon.

We note that this price-based adaptive terminal cost addresses only one component of the longer-horizon planning problem. This is clearly a limitation of the terminal costification and hence of the proposed control mechanism. In principle, the optimal terminal penalty depends not only on future electricity prices but also on anticipated travel demand, the fleet-wide state-of-charge distribution, spatial vehicle positioning, and other state-dependent factors. This limitation reflects a fundamental challenge in receding horizon control: the terminal value function should ideally approximate the optimal cost-to-go, which depends on the full system state and is itself the solution to an optimization problem of the same dimensionality as the original planning problem (Bertsekas, 2024; Rawlings et al., 2017). Our price-based formulation addresses only the energy cost component of this value function, leaving demand-driven and fleet-state-driven components unmodeled. The formulation in Eq. (11) should therefore be understood as a tractable heuristic that encodes readily available day-ahead price information, rather than a comprehensive solution to the time-scale mismatch.

3.4. The multi-stage SMPC problem

The EAMoD control problem for rebalancing and charging under uncertainty can now be formulated using the model and objective function introduced above. The overall goal is to improve both service quality and operational efficiency. This is achieved by strategically rebalancing vehicles to match spatial demand patterns and by scheduling charging and V2G operations to maintain adequate battery SoC.

Uncertainty in the future evolution of the random variables is addressed using a combined scenario-based and robust approach. Travel demand and charger availability are modeled by a scenario tree that approximates the expected value, with each realization indexed by the superscript s . Uncertainties in energy consumption and travel time are handled through the chance-constrained parameterization of the model parameters Δb_{ijt} , Δt_{jit} , and $b_{ij}^{\min}$ described in Section 3.2.1, which ensures operational feasibility with high probability. The resulting optimization problem is updated sequentially as new information about the random variables becomes

available. The multi-stage SMPC problem is given by

$$\min_{x_{ijbt}^{r,s}, x_{ijbt}^{p,s}, x_{iibt}^{c,s}, x_{iibt}^{d,s}} J_N(x_{ijbt}^{r,s}, x_{ijbt}^{p,s}, x_{iibt}^{c,s}, x_{iibt}^{d,s}) \quad (12a)$$

subject to

$$(1), (4), (6), \quad (12b)$$

$$x_{ijbt}^{r,s} = x_{ijbt}^{r,s'}, \quad x_{ijbt}^{p,s} = x_{ijbt}^{p,s'}, \quad x_{iibt}^{c,s} = x_{iibt}^{c,s'}, \quad x_{iibt}^{d,s} = x_{iibt}^{d,s'} \quad \forall s, s' \in S_t : (\xi_0^s, \dots, \xi_{t-1}^s) = (\xi_0^{s'}, \dots, \xi_{t-1}^{s'}), \quad (12c)$$

$$x_{ijbt}^p, x_{iibt}^c, x_{iibt}^d, s_{ijt} \in \mathbb{N} \quad \forall i, j \in \mathcal{N}, t \in \mathcal{T}, b \in B, \quad (12d)$$

$$x_{ijbt}^r \in \mathbb{R}^+ \quad \forall i, j \in \mathcal{N}, t \in \mathcal{T}, b \in B. \quad (12e)$$

Constraint (12c) enforces non-anticipativity: at each stage, decisions must be identical across all scenarios that share the same history. Note that the uncertainties in energy consumption and travel time do not appear as explicit constraints in the optimization. Instead, they are handled through the chance-constrained parameterization of the model described in Section 3.2.1: the conservative travel time estimates $\hat{\tau}_{jit}$ from (2) enter the flow conservation constraint (1), and the minimum SoC threshold $b_{ij}^{\min}$ from (5) restricts vehicle dispatch in the imbalance constraint (4). These parameters are computed offline at each receding horizon iteration before the optimization is solved, effectively embedding the robust treatment of energy and travel time uncertainty into the model structure itself.

The problem is nearly totally unimodular (TU). In general, the LP relaxation of a TU mixed-integer linear program (MILP) yields integer solutions (Hoffman and Kruskal, 2010). In our case, however, two constraints break TU. First, Eq. (4) enforces integrality on the sum of passenger-serving variables x_{ijbt}^p rather than on each variable individually. Second, Eq. (6) couples charging and discharging decisions across vehicles in a way that also violates TU. Consequently, explicit integrality must be imposed on x_{ijbt}^p, x_{iibt}^c , and x_{iibt}^d . Only the rebalancing variables x_{ijbt}^r can be treated as positive continuous flows. The overall optimization problem therefore remains a MILP but cannot exploit full TU structure. A significant computational challenge in solving Eq. (12) arises from the exponential growth in the number of scenarios considered at each stage, as also noted by Richards and How (2005). To address this, the next section introduces an integrated scenario prediction and reduction method, together with a decomposition technique. The resulting extensive-form MILP grows exponentially with the number of stages and scenarios, rendering direct solution impractical for real-time MPC applications.

4. Scenario tree generation and reduction

To optimize fleet rebalancing and charging under uncertainty, we use a SMPC framework. A key challenge is the uncertain evolution of future travel demand and charger availability, which prevents a closed-form evaluation of the expected cost. We address this by representing uncertainty with a scenario tree that captures both multiple realizations and temporal dependence. This section outlines the construction of the scenario tree, the Bayesian Neural Network (BNN)-based forecasting methodology, the scenario reduction procedure, and the integration of the reduced tree into the SMPC formulation.

4.1. Scenario tree construction

The true distributions D_λ^* and D_k^* for travel demand and charger availability are unknown, so the expectation in (8) cannot be computed directly. To approximate this, we represent uncertainty using a scenario tree, which encodes multiple realizations of the random variables and their evolution over time (see Fig. 3). At each time step t , a scenario $s \in \{1, 2, \dots, S\}$ is defined as $\xi_t^s = [\lambda_{ijt}^s, k_{it}^{c,s}]$, where λ_{ijt}^s is the predicted travel demand and $k_{it}^{c,s}$ the available chargers in scenario s . The cost function is then reformulated using the Sample Average Approximation (SAA) approach:

$$\begin{aligned} \hat{J}_N(x_{ijbt}^{r,s}, x_{ijbt}^{p,s}, x_{iibt}^{c,s}, x_{iibt}^{d,s}) &= \sum_{i=1}^N \sum_{j=1}^N \sum_{b=0}^1 Q(x_{ijb0}^r, x_{ijb0}^p, x_{iib0}^c, x_{iib0}^d) \\ &+ \sum_{i=1}^{T-1} \sum_{s=1}^{S_i} p_s \sum_{i=1}^N \sum_{j=1}^N \sum_{b=0}^1 Q(x_{ijbt}^{r,s}, x_{ijbt}^{p,s}, x_{iibt}^{c,s}, x_{iibt}^{d,s}) \\ &+ \sum_{s=1}^{S_T} p_s \sum_{i=1}^N \sum_{j=1}^N \sum_{b=0}^1 C^{\text{term}}(b) \cdot (x_{ijbT}^{r,s} + x_{ijbT}^{p,s} + x_{iibT}^{c,s} + x_{iibT}^{d,s}), \end{aligned} \quad (13)$$

where p_s denotes the probability of scenario s , and $\sum_{s=1}^S p_s = 1$. The function $\hat{J}_N(\cdot)$ approximates the original objective in Eq. (8), where the expected value is replaced with a scenario tree. To avoid exponential growth of the scenario tree with respect to the prediction horizon, we introduce the concept of a *robust horizon*, beyond which uncertainty is assumed to remain constant. This idea, proposed in Lucia et al. (2013), allows branching only up to a certain stage, making the tree more tractable (see Fig. 3). Additionally, to further limit unnecessary tree growth when the time discretization is fine, scenarios are assumed to remain constant over short time intervals, effectively reducing the number of unique branches in closely spaced time steps.




Fig. 3. Scenario tree representation of the random variables, incorporating both the robust horizon and prediction horizon. Each scenario is defined as $\xi_t^s = (\lambda_{ijt}^s, k_{it}^{c,s})$.

4.2. Forecasting travel demand and charger availability

Given historical spatio-temporal data on travel demand and charger availability, our goal is to predict future values, quantify uncertainty, and model correlations across time and space. Input features x include time, location, and contextual variables such as weather or day of the week. Travel demand is recorded as the number of trips between discretized spatial zones over fixed time intervals, resulting in matrices $\lambda_{ijt} \in \mathbb{Z}^{N \times N}$. Charger availability is given as vectors $k_{it}^c \in \mathbb{Z}^N$. Because both outputs are non-negative counts, the predictive model must produce integer-valued outputs and capture predictive uncertainty.

We use a Bayesian Neural Network (BNN) to model these forecasts. Unlike standard neural networks, BNNs maintain distributions over their weights θ , allowing for explicit modeling of epistemic uncertainty. Given input features x , the neural network defines a parametric function $f(x; \theta)$ that produces an output $\hat{y}$. The Bayesian approach introduces a posterior distribution over the parameters:

$$p(\theta|D) \propto p(D|\theta)p(\theta), \quad (14)$$

where $p(\theta)$ is the prior over weights and $p(D|\theta)$ is the likelihood of the observed data given parameters θ . Prediction for a new input x involves marginalizing over this posterior:

$$p(\hat{y}|x, D) = \int p(\hat{y}|x, \theta)p(\theta|D)d\theta. \quad (15)$$

This integral is intractable in general, so approximate inference techniques such as Variational Inference (VI) or Markov Chain Monte Carlo (MCMC) are used. In our work, we adopt the BNF model proposed in [Saad et al. \(2024\)](#), which is specifically designed for spatio-temporal prediction with multiple correlated outputs. It models the predictive distribution using the Negative Binomial (NB) distribution, which is well-suited for overdispersed count data:

$$p(\hat{y}(x)|D) \sim \text{NB}(\hat{\mu}, \hat{\sigma}^2),$$

where $\hat{\mu}$ and $\hat{\sigma}^2$ are the mean and variance predicted by the model. These predictive distributions form the basis for scenario generation in the next step of our stochastic optimization framework. An important advantage of the Bayesian formulation is that the model can be updated sequentially as new data arrive. In the variational inference setting, this is achieved by using the previously learned posterior as the prior when incorporating a new dataset D_t . This enables efficient online learning without retraining from scratch, while preserving past information. Such posterior-to-prior updates allow the BNN to adapt to changing travel demand patterns or charger availability in real time, which is crucial for EAMoD operations. These predictive distributions, updated over time, form the basis for scenario generation in the next step of our stochastic optimization framework.

The following summarizes the key hyperparameters of the BNF model used in our experiments. A daily seasonality period with 7 harmonics captures intra-day demand patterns such as morning and evening peaks. The feature set consists of five columns: the time interval index and the spatial coordinates of both pickup (x, y) and dropoff (x_d, y_d) locations. To capture joint spatio-temporal effects, all $\binom{5}{2} = 10$ pairwise feature interactions are included as explicit second-order cross-product inputs to the network. These interactions fall into three categories: (i) spatiotemporal coupling (e.g., time $\times$ pickup longitude), which allows demand at different locations to vary independently over the day; (ii) origin–destination flow structure (e.g., pickup longitude $\times$ dropoff longitude), which captures directional travel patterns between regions; and (iii) neighborhood identification (e.g., pickup $x \times$ pickup y), which enables the model to distinguish specific geographic areas. By providing these cross-products explicitly, the model can learn multiplicative relationships efficiently even with a compact architecture (width 64, depth 2), rather than relying solely on the hidden layers to discover them.

4.3. Scenario tree reduction

From the BNF model, we generate scenario samples of the form:

$$\xi^s = (\xi_1^s, \xi_2^s, \dots, \xi_T^s), \quad s = 1, \dots, S,$$

where each ξ_t^s consists of sampled travel demand and charger availability values for scenario s at time t . This collection of sampled sequences is referred to as a *scenario fan*. To ensure tractability, we reduce the number of scenarios using clustering techniques. One common approach is to treat the problem as a two-stage stochastic program, but this requires a large number of samples. Instead, we construct a reduced scenario tree that preserves the structure of a multi-stage stochastic process. The tree is built iteratively from the root. At each stage, we sample future values conditionally using the predictive model:

$$\xi_{(t+1)}^s \sim P(\hat{y} \mid D, \hat{y}(x(t))).$$

We then apply soft k -means clustering to reduce the number of scenarios, using the cluster centroids as representatives for each scenario. This process is repeated for each time stage. The number of clusters K_t at each step is specified by the user in order to balance tractability and representativeness of the uncertainty.

Alternative reduction approaches based on probability metrics, such as the Wasserstein distance or the Nested distance, are not adopted here due to their high computational cost and limited scalability in multi-stage settings. Instead, our method follows the general idea of clustering-based reductions that have been proposed in previous works, where scenarios are grouped and replaced by representative centroids. This allows us to construct reduced scenario trees that remain computationally feasible while still capturing the main patterns of uncertainty. The full reduction algorithm is given in [Algorithm 1](#).

Algorithm 1 Scenario tree reduction.

- 1: **Input:** Initial state ξ_0 , dataset D , prediction model $P(\cdot \mid D)$, time horizon T , number of scenarios per stage K_t
 - 2: **Output:** Reduced scenario tree
 - 3: Initialize root node with ξ_0
 - 4: **for** $t = 0$ to $T - 1$ **do**
 - 5: Sample M scenarios: $\xi^s(t+1) \sim P(\hat{y}(x) \mid D, \hat{y}(x(t)))$
 - 6: Apply soft k -means to obtain K_t representative scenarios (with K_t specified by the user)
 - 7: Set cluster centroids
 - 8: Assign transition probabilities from clustering weights
 - 9: **end for**
 - 10: **return** Constructed scenario tree
-

5. Nested Benders decomposition

Directly solving the multi-stage stochastic problem in [Eq. \(12\)](#) can be computationally expensive, and even intractable, due to the exponential growth of scenarios over the planning horizon. To address this challenge, we employ Nested Benders Decomposition (NBD), an extension of classical Benders decomposition tailored for multi-stage stochastic programs. The key idea is to decompose the original problem into a hierarchy of two-stage problems, one for each stage and scenario, while iteratively exchanging information between them (see [Figs. 4 and 5](#)). This approach reduces the computational burden while preserving anticipativity across the scenario tree.

In the NBD framework, the problem is reformulated as a sequence of nested two-stage subproblems, each associated with a specific stage and scenario ([Fig. 4](#)). A master problem (MP) coordinates the solution process by integrating Benders cuts derived from the duals of the subproblems (SP) ([Fig. 5](#)). These cuts iteratively refine the feasible region of the MP, guiding the solution process toward global optimality while reducing computational effort.

The NBD procedure is executed in a forward-backward iterative scheme, commonly referred to as fast-forward-fast-back sequencing ([Murphy, 2013](#)). During the forward pass, the first-stage problem and subsequent stages, up to one stage before the robust horizon, are treated as master problems. In the backward pass, only the first stage remains a master problem, while all subsequent stages are treated as subproblems. To compute valid duals, integer variables in these subproblems must be relaxed to continuous ones.

The iterative process terminates once convergence is achieved, defined as the gap between the subproblem costs and their estimated values in the MP falling below a specified tolerance, or when a maximum number of iterations is reached. Because subproblem integer variables are relaxed, the resulting solutions may not be integer-feasible. This relaxation introduces a *duality gap* between the relaxed subproblem objective and the true integer-optimal objective, which weakens the tightness of the Benders cuts. Techniques such as Gomory cuts ([Gade et al., 2014](#)) and Lagrangian cuts ([Chen and Luedtke, 2022](#)) can yield tighter cuts, but they often significantly slow convergence and are therefore not adopted in our implementation.

Within the NBD framework, multiple MPs are defined, and as the decision process progresses down the scenario tree, SPs are recursively promoted to MPs. To formalize this structure, we present general formulations for both MPs and SPs. Based on the SMPC formulation in [Eq. \(12\)](#) and the reformulated objective function in [Eq. \(13\)](#), the first-stage master problem can be expressed as follows:

$$\text{minimize}_{x_{ijbt}^{r,s}, x_{ijbt}^{p,s}, x_{iibt}^{c,s}, x_{iibt}^{d,s}} J_N(x_{ijbt}^{r,s}, x_{ijbt}^{p,s}, x_{iibt}^{c,s}, x_{iibt}^{d,s}) + \sum_{s=1}^S p_s \alpha_s \quad (16a)$$

subject to

$$(1), (4), (6), (12c) \quad (16b)$$




Fig. 4. An example of solving a multi-stage SMPC problem by decomposing it into a nested series of two-stage SMPC problems. The figure on the left illustrates the first and second stages of the SMPC problem, while the figure on the right shows a portion of the second and third stages. Both sections can be resolved using Benders Decomposition.




Fig. 5. Illustration of nested Benders decomposition. The master problem proposes trial solutions which are passed sequentially to subproblems at different stages. Each subproblem generates Benders cuts based on its solution and returns them to the previous stage, enabling iterative refinement across the scenario tree.

$$\alpha_s \geq Q^s + \sum_{ijb} \pi_{ijbt}^{r,s} (x_{ijbt}^{r,s} - \bar{x}_{ijbt}^{r,s}) + \pi_{ijbt}^{p,s} (x_{ijbt}^{p,s} - \bar{x}_{ijbt}^{p,s}) + \pi_{iibt}^{c,s} (x_{iibt}^{c,s} - \bar{x}_{iibt}^{c,s}) + \pi_{iibt}^{d,s} (x_{iibt}^{d,s} - \bar{x}_{iibt}^{d,s}) \quad (16c)$$

$$\forall s \in S_{t+1} \quad (16d)$$

$$x_{ijbt}^p \in \mathbb{N} \quad \forall i, j \in N, t \in \mathcal{T}, \forall b \in \mathcal{B}. \quad (16e)$$

$$x_{ijbt}^r, x_{iibt}^c, x_{iibt}^d \in \mathbb{R}^+ \quad \forall i, j \in \mathcal{N}, \forall t \in \mathcal{T}, \forall b \in \mathcal{B}, \quad (16f)$$

where p_s is the probability of each scenario, S_t is the number of scenarios in the stage, and α_s is the approximated lower bound of the cost of the subproblem. The MP inherits all structural constraints from the SMPC formulation (12), including flow conservation, imbalance dynamics, charger capacity, and non-anticipativity. The chance constraints on energy consumption and travel time are embedded in the model parameters (see Section 3.2.1) and thus implicitly enforced through (1) and (4). These constraints are augmented by the Benders cuts in constraint (16c), which approximate the future cost. Constraint (16c) represents the Benders cuts from the subproblems at the next time step, and S_{t+1} denotes the set of all subproblems at time $t + 1$. This version of the Benders cuts is referred to as the compact form. In this constraint, Q^s represents the cost of each next-stage subproblem, and π denotes the dual variables associated with the corresponding decision variables in the MP. The BD cuts generated in these MPs will always be optimality cuts due to the flow conservation constraints in Eq. (1). As a result, feasibility cuts are not required. In this implementation, we use multi-cuts, where one Benders cut is added to the master problem for each subproblem. Alternatively, a single-cut approach can be used, where the cuts from each subproblem are aggregated into a single cut. The advantages and disadvantages of both approaches are discussed in Chapter 5 of Birge and Louveaux (2011). It should be noted that whenever an optimization problem is a master problem, integrality is enforced on $x_{ijbt}^{p,s}$. However, when moving up the scenario tree and the MP becomes a subproblem, the problem is relaxed to a linear program to enable the computation of dual variables. The subproblem is then,

$$\text{minimize}_{x_{ijbt}^{r,s}, x_{ijbt}^{p,s}, x_{iibt}^{c,s}, x_{iibt}^{d,s}} J_N(x_{ijbt}^{r,s}, x_{ijbt}^{p,s}, x_{iibt}^{c,s}, x_{iibt}^{d,s}) + \sum_{s=1}^S p_s \alpha_s \quad (17a)$$

subject to

$$(1), (4), (6), (12c) \quad (17b)$$

$$\alpha_s \geq Q^s + \sum_{ijb} \pi_{ijb}^{r,s} (x_{ijb}^{r,s} - \bar{x}_{ijb}^{r,s}) + \pi_{ijb}^{p,s} (x_{ijb}^{p,s} - \bar{x}_{ijb}^{p,s}) + \pi_{iib}^{c,s} (x_{iib}^{c,s} - \bar{x}_{iib}^{c,s}) + \pi_{iib}^{d,s} (x_{iib}^{d,s} - \bar{x}_{iib}^{d,s}) \quad (17c)$$

$$\forall s \in S_{t+1} \quad (17d)$$

$$x_{ji(b+\Delta b_{ji})(t-\Delta t_{ji})}^p = \bar{x}_{ji(b+\Delta b_{ji})(t-\Delta t_{ji})}^p \quad (17e)$$

$$x_{ji(b+\Delta b_{ji})(t-\Delta t_{ji})}^r = \bar{x}_{ji(b+\Delta b_{ji})(t-\Delta t_{ji})}^r \quad (17f)$$

$$x_{ii(b-\Delta b^c)(t-1)}^c = \bar{x}_{ii(b-\Delta b^c)(t-1)}^c \quad (17g)$$

$$x_{ii(b+\Delta d)(t-1)}^d = \bar{x}_{ii(b+\Delta d)(t-1)}^d \quad (17h)$$

$$s_{ij(t-1)} = \bar{s}_{ij(t-1)} \quad (17i)$$

$$x_{ijbt}^p, x_{ijbt}^r, x_{iibt}^c, x_{iibt}^d \in \mathbb{R}^+ \quad \forall i, j \in \mathcal{N}, \forall t \in \mathcal{T}, \forall b \in \mathcal{B}. \quad (17j)$$

where constraints (17e)–(17i) are the solution from the previous time step MP. The SP have the BD cuts from previous iteration of the NB, (17c).

5.1. Improving nested benders decomposition

To accelerate the convergence of the NBD algorithm, we employ several enhancement techniques. First, subproblems corresponding to the same stage can be solved in parallel, which can significantly reduce computational time and particularly when stage aggregation is used (Dempster and Thompson, 1998). Aggregating multiple stages reduces the depth of the decomposition tree and shifts computational effort from cut generation to problem solving. While solving each aggregated subproblem becomes more time-consuming, the overall number of Benders cuts required typically decreases, leading to faster convergence.

So-called warm-up cuts are implemented prior to starting the NBD algorithm. These cuts are based on simple heuristics. For each scenario, the vehicle imbalance at every station and time step is computed in the absence of rebalancing. Based on these imbalances, slack-penalizing constraints are added to the master problems to incentivize early rebalancing and charging actions. The cost of the slack variable corresponds to the lost revenue from not serving a customer. These warm-up cuts help guide the initial iterations toward more balanced solutions and can significantly reduce convergence time.

6. Mobility network simulations and sensitivity analysis

We evaluate our approach in two high-fidelity case studies, San Francisco and Chicago, constructed from large-scale taxi trip records. We first assess the accuracy of the proposed travel-demand predictions on these data. Next, we analyze the computational performance of the Nested Benders Decomposition (NBD) used to solve the underlying stochastic optimization problems. We then benchmark the proposed stochastic and robust algorithm in the AMoDeus simulator (Ruch et al., 2018). Finally, we perform a sensitivity analysis to quantify how vehicle energy efficiency, battery capacity, fleet size, and charger availability affect system performance.

6.1. Simulation environment

We evaluate our proposed algorithm using the high-fidelity transport simulator AMoDeus (Ruch et al., 2018), an extension of the Multi-Agent Transport Simulation framework (MATSim) (Horni et al., 2016) specifically designed for AMoD systems. Two urban contexts are studied: San Francisco and Chicago, both based on real-world taxi data.

The San Francisco case study utilizes taxi data from 2008, encompassing 464,045 customer trips collected from 500 taxis (Piorkowski et al., 2022). For simulation, we selected Thursday, May 29, 2008 with a total travel demand of 11,453 trips and an average trip distance of 3.31 km. The Chicago scenario is based on taxi data collected between 2013 and 2023, representing over 212 million trips (City of Chicago, 2023). Our analysis focuses on Wednesday, July 24, 2019 which recorded a total travel demand of 32,406 trips and an average trip distance of 14.64 km, using data from July 15–24, 2019. To meet the respective demand levels, we simulate fleets of 500 vehicles in San Francisco and 1600 in Chicago.

In both cities, vehicles are initially distributed across $\mathcal{N} = 10$ charging stations, each starting at 80% SoC. The SoC is discretized into $B = 50$ intervals. The simulation horizon is $T = 20$ time steps (200 minutes) with a time step length of $\Delta t = 10$ minutes, and operational planning is carried out over a robust horizon of 6 time steps (60 minutes). Vehicles are required to end the simulation with at least $b_{\text{end}} = 30\%$ SoC. An overview of all simulation parameters is provided in Table 3.

Vehicle battery capacities are set to $e^b = 40$ kWh for San Francisco and $e^b = 70$ kWh for Chicago. We use a non-linear charging curve derived from Wassiliadis et al. (2021), in which charging proceeds at a constant 100 kW until 45% SoC and then decreases linearly to 30 kW at 95% SoC. The stochastic optimization model applies scenario reduction per time step with levels [3, 3, 5, 5, 5], resulting in 1125 total scenarios. Battery violation tolerance is set to $\epsilon^b = 0.2$ and demand violation tolerance to $\epsilon^r = 0.2$. Nested Benders Decomposition (NBD) is used with a gap tolerance of 10^{-2} and a maximum of 30 iterations. The terminal cost parameter a_t is updated at each receding horizon iteration based on day-ahead electricity prices from Nordpool (Nord Pool, 2025), using a lookahead window of $H = 36$ time steps (6 hours) and a baseline parameter $a_{\text{base}} = 0.10$ €/kWh.

A key enhancement to the simulation environment is the integration of an energy consumption estimation module, as AMoDeus does not include this functionality for transportation simulations. This module estimates energy consumption using detailed speed profiles based on the established model from (Basso et al., 2021). Furthermore, to accurately represent real-world driving conditions, stochasticity was introduced into the simulation by modeling traffic light behavior and randomizing vehicle speeds.




Fig. 6. Electricity prices in euro/kWh for three different pricing scenarios: flat, peak, and solar price, shown over a 24-hour period.

Table 3

Simulation parameters overview.

Common Parameters		
Parameter	Value	
Number of stations (N)	10	
Simulation horizon (T)	20 time steps (200 minutes)	
Robust horizon	6 time steps (60 minutes)	
Scenarios reduction per time step	[3,3,5,5,5] (1125 scenarios in total)	
Time step length (Δt)	10 minutes	
Number of SoC intervals (B)	50	
Minimum required SoC at end (b_{end})	30%	
V2G power (P^d)	22 kW	
Battery violation tolerance (ϵ^b)	0.2	
Demand violation tolerance (ϵ^d)	0.2	
Vehicle type	V1	
NBD gap tolerance	10^{-2}	
Max number of BD iteration	30	
City-Specific Parameters		
Parameter	San Francisco	Chicago
Number of vehicles	500	1600
Battery capacity (e^b)	40 kWh	70 kWh
Total number of chargers	100	200
Total demand	11,453	32,406
Simulation date	May 29, 2008	July 24, 2019

For electricity pricing, we developed three distinct scenarios (Fig. 6): a flat-price profile, a peak-price profile, and a solar-price profile reflecting midday price drops from high solar generation, with the first two derived from Nordpool data (Nord Pool, 2025). These represent daily operational variations rather than annual averages. For benchmarking the proposed algorithm, vehicle type vehicle 1 is used, while the sensitivity analysis compares three different vehicle types with varying energy consumption rates (Fig. 8), under the peak-price scenario only.

All simulations were executed on a MacBook Pro with a 2.3 GHz Quad-Core Intel Core i7 processor and 16 GB of RAM. Optimization problems were solved using the Gurobi optimization library (Gurobi Optimization, 2022).

6.2. Travel demand prediction

We conducted a comprehensive comparative analysis between BNF, our proposed prediction framework, and three established baselines: multi-output Gaussian Process regression with a spectral kernel (MOSK) (Parra and Tobar, 2017), Long Short-Term Memory (LSTM) networks, and Vector AutoRegressive Moving Average (VARIMA) models (Zivot and Wang, 2006). These baselines were selected based on their proven effectiveness in travel demand forecasting and represent a diverse set of statistical and machine learning approaches (Iglesias et al., 2018; He et al., 2023). Like BNF, MOSK provides explicit uncertainty estimates, enabling a direct comparison of probabilistic forecasting performance. All models were evaluated on real-world datasets from San Francisco and Chicago. The BNF architecture used two hidden layers with 64 units each. Input features included time, origin coordinates (latitude and longitude), and destination coordinates. A Negative Binomial distribution was used as the prior predictive model to reflect the count nature of the data.

Table 4

Comparison of RMSE and MIS for different forecasting models in San Francisco and Chicago across key time periods.

Metric	Algorithm	San Francisco			Chicago		
		06:00–10:00	10:00–14:00	17:00–21:00	06:00–10:00	10:00–14:00	17:00–21:00
RMSE	LSTM	14.5175	18.8251	24.2956	27.4708	85.3836	133.4684
	VARIMA	18.0840	9.6063	12.0309	29.3161	26.6069	35.5457
	MOSK	10.0874	7.8890	10.5079	26.0935	26.0016	28.8127
	BNF	8.1455	6.9751	8.0259	25.2362	13.6507	14.6357
MIS	MOSK	83.2670	63.8836	77.2634	193.0978	154.8194	148.6521
	BNF	40.0050	33.6349	53.0234	110.9026	98.2883	103.8081




Fig. 7. Predicted travel demand with 95% prediction intervals for selected locations. Black dots show training data, white dots indicate test data, the red line is the median prediction, and blue shaded areas represent the 95% prediction intervals.

Model performance was assessed using two key metrics which were Root Mean Square Error (RMSE) and Mean Interval Score (MIS). RMSE quantifies the average deviation between predicted and actual values, with lower values indicating higher accuracy. MIS penalizes both the width and coverage of prediction intervals, providing a measure of uncertainty calibration; lower values correspond to sharper and more reliable probabilistic forecasts. Results are shown in Table 4.

BNF achieved the best performance across all scenarios in both RMSE and MIS. In San Francisco, BNF reduced RMSE significantly by an average of 57.9% compared to LSTM, 38.6% compared to VARIMA, and 18.1% compared to MOSK. The improvements were more pronounced in Chicago, with average RMSE reductions of 60.4%, 40.5%, and 33.3% compared to LSTM, VARIMA, and MOSK, respectively. For uncertainty quantification, BNF also demonstrated a clear advantage, reducing MIS by an average of 43.6% in San Francisco and 36.4% in Chicago relative to MOSK.

In addition to demonstrating strong average performance, BNF showed greater robustness during peak demand periods (06:00–10:00 and 17:00–21:00), which are especially important for real-world EAMoD operations. During these high-demand windows, models such as MOSK exhibited significantly higher errors. For example, during the evening peak in Chicago, BNF reduced RMSE by 49.2% compared to MOSK, indicating its superior ability to handle sharp fluctuations in travel demand. The substantial reductions in both RMSE and MIS contribute to more reliable scenario trees. A lower RMSE results in a more accurate average scenario, while a lower MIS improves the estimation of scenario probabilities within the scenario tree.

As shown in Fig. 7, BNF generates accurate travel demand forecasts with well-calibrated 95% prediction intervals across selected Chicago pickup locations. While the median prediction does not perfectly match realized demand, the uncertainty estimates provide valuable additional information. This highlights the importance of considering predictive uncertainty in downstream decision-making and motivates the construction of scenario trees based on the probabilistic forecasts.

6.3. Nested Benders decomposition performance

We analyze the computational performance of Nested Benders Decomposition (NBD), which is used to solve the multistage stochastic optimization problems underlying our SMP approach. We do not compare against the extensive form of the stochastic problem, as it becomes computationally intractable beyond very small instances. Instead, we evaluate NBD in terms of runtime, convergence behavior, and scalability with respect to the number of scenarios.

In the nested Benders decomposition, one optimality gap is computed for each master problem and subproblem across the scenario tree. Since these gaps can vary considerably across nodes, we report both the maximum and the median gap as indicators of convergence quality. The maximum gap highlights the worst-case convergence behavior across all nodes, while the median gap reflects the typical convergence performance and is less sensitive to outliers. The optimality gap reported here refers specifically to the Benders

Table 5

Runtime, number of iterations, and relative optimality gaps of NBD across different scenario counts.

Scenarios	Runtime (sec)	Iterations	Max Gap (%)	Median Gap (%)
135	17.15	10	0.08	0.08
540	90.26	9	0.79	0.09
1125	355.55	15	4.52	1.50

decomposition procedures. In each iteration, the Benders cuts provide a lower bound on the subproblem cost for the current master solution, while the incumbent solution provides an upper bound. The relative gap is measured as

$$\text{Gap} = \frac{UB - LB}{UB} \times 100\%,$$

where UB denotes the incumbent upper bound and LB the lower bound from the cuts. In the ideal case, this gap equals 0%, which means that the decomposition have reached the exact optimum. However, driving the gap fully to zero can be computationally prohibitive. In practice, we terminate the NBD algorithm if the gap does not improve for five consecutive iterations or if the runtime exceeds 600 seconds. The reported gaps reflect only the internal convergence of the decomposition. Since the subproblems are solved in relaxed form, the true duality gap with respect to the extensive formulation remains unobservable without solving the full problem.

Table 5 reports the runtime, number of iterations, and the relative optimality gaps across three scenario set sizes. For the smallest instance (135 scenarios), NBD converges within 10 iterations and achieves a near-exact solution with a final max gap of 0.08%. Increasing the number of scenarios to 540 raises the runtime to about 90 seconds, but convergence is still reached in fewer than 10 iterations, with a small median gap of 0.09%.

For the largest instance (1125 scenarios), the runtime increases to nearly 6 minutes and the number of iterations to 15. While the maximum gap observed during the iterations is higher (4.52%), the algorithm ultimately converges to a high-quality solution with a median gap of 1.50%. These results demonstrate that, although runtime grows with the number of scenarios, NBD remains computationally tractable and consistently delivers solutions with small optimality gaps, even for problem sizes where solving the extensive form is impractical.

6.3.1. Sensitivity to scenario tree configuration

To evaluate the impact of scenario tree design on operational performance, we conducted a sensitivity analysis varying both the total number of scenarios and their allocation across stages. We compared two allocation strategies with equal total scenarios: (i) early branching [5,5,5,3,3], which allocates more scenarios to early stages to resolve near-term uncertainty; and (ii) late [3,3,5,5,5], our baseline configuration, which concentrates branching in the later stages. Additionally, a medium configuration [3,3,3,4,5] with 540 total scenarios and a small configuration [3,3,3,5] with 135 total scenarios are included to assess the effect of decreasing the total scenario count. Table 6 reports the results for San Francisco and Chicago. The results show a modest but consistent improvement in tail performance as the total number of scenarios increases, while median wait times remain largely stable across configurations. In San Francisco, moving from the Small to the Large configuration leaves the median wait time unchanged at 229 s, while the 95th percentile decreases from 559 s to 548 s. In Chicago, the median similarly remains at 214 s, with the 95th percentile reducing from 783 s to 767 s, reflecting the higher sensitivity to uncertainty resolution under greater demand intensity. Beyond 1125 scenarios, gains plateau while computational cost increases substantially, justifying this as an appropriate upper bound. For a fixed total of 1125 scenarios, the Late [3,3,5,5,5] configuration achieves a marginally lower 95th percentile wait time than the Early [5,5,5,3,3] configuration in both cities (548 s vs. 550 s in San Francisco; 767 s vs. 767 s in Chicago), with comparable total and rebalancing distances. This can be attributed to the heterogeneous nature of forecast uncertainty across the planning horizon: as shown in Fig. 7, prediction uncertainty is relatively low in the near-term stages and increases toward later stages. Consequently, allocating higher branching factors to early stages, where scenarios are nearly indistinguishable, yields little additional representational value, whereas denser branching in later stages better captures the growing spread in demand and charger availability. On this basis, we adopt the Late [3,3,5,5,5] configuration as the default, as concentrating branching in stages with greater uncertainty leads to a more informative scenario tree for a given computational budget.

6.4. EAMoD optimization performance comparison

We benchmark our proposed SMPC algorithm, which explicitly accounts for uncertainty in predictions, against three alternative control strategies:

- **Deterministic Mean Prediction Optimization (DMP):** Uses only the expected values of predicted travel demand, charging station occupancy, travel time, and energy consumption from BNF, without considering uncertainty.
- **Robust Chance-Constrained Optimization (RCC):** A robust approach inspired by Jacobsen et al. (2023), designed to guarantee performance under worst-case prediction errors.
- **Reactive Baseline (RB):** A reactive strategy with no forecasting. Vehicles are charged whenever the state of charge (SoC) falls below 20%, up to 80%. No vehicle-to-grid (V2G) operations are considered.

Table 6

Sensitivity analysis of scenario tree configuration for San Francisco (fleet size 500) and Chicago (fleet size 1600). Metrics include median and 95th percentile wait times (95% WT), total distance and rebalancing distance per vehicle.

City	Configuration	Scenarios (total)	Median WT [s]	95% WT [s]	Total Dist./Veh. [km]	Reb. Dist./Veh. [km]
Chicago	Small [3, 3, 3, 5]	135	214	783	495.2	69.1
	Medium [3, 3, 3, 4, 5]	540	214	767	500.7	70.1
	Large [3, 3, 5, 5, 5]	1125	214	767	500.7	70.1
	Early [5, 5, 5, 3, 3]	1125	214	767	506.3	69.8
San Francisco	Small [3, 3, 3, 5]	135	229	559	412.4	103.7
	Medium [3, 3, 3, 4, 5]	540	229	555	413.2	106.3
	Large [3, 3, 5, 5, 5]	1125	229	548	414.1	108.5
	Early [5, 5, 5, 3, 3]	1125	229	550	412.5	107.5

All controllers are re-optimized every 10 minutes using a receding horizon approach. The results for Chicago and San Francisco are summarized in [Tables 8](#) and [9](#).

In Chicago, SMPC achieves the lowest passenger delays. The median waiting time is 214 s, essentially the same as for DMP (215 s), but significantly lower at the 95th percentile with 767 s compared to 872 s for DMP (12% improvement), 822 s for RCC, and 955 s for the reactive baseline (20% improvement). Operationally, SMPC requires 70.1 km of rebalancing per vehicle, which is 27% less than RCC (96.6 km), while keeping the total distance traveled (500.7 km) close to RCC (510.1 km). RB travels the least overall distance (456.1 km), but this comes at the cost of much longer wait times.

In San Francisco, SMPC again dominates in terms of service quality. The median waiting time is 229 s, compared to 240 s for DMP (5% higher), 254 s for RCC, and 360 s for the reactive baseline, corresponding to a 36% improvement over RB. At the 95th percentile, SMPC records 548 s compared to 584 s for DMP, 559 s for RCC, and 670 s for RB, improving extreme waits by almost 20% relative to the reactive baseline. Operationally, SMPC requires 108.5 km of rebalancing per vehicle compared to 124.8 km for RCC, a 13% reduction, while keeping total distance traveled (414.1 km) lower than RCC (431.6 km). Although DMP uses slightly less rebalancing (99.5 km), SMPC achieves consistently lower waiting times across the distribution, showing that the modest increase in repositioning yields substantially improved service reliability.

Energy costs and V2G revenues show similar trends in both cities. Under flat tariffs, all methods incur 0.06 euro/kWh in costs and generate no revenue, as expected. However, under peak and solar tariffs, SMPC and DMP consistently achieve the lowest effective energy costs and the highest revenues. In Chicago, SMPC earns almost 975 euro under solar pricing, which is approximately double the value in San Francisco of 509 euro, reflecting the larger fleet and heavier charging activity. RCC captures somewhat less revenue due to its conservative scheduling, while RB never participates in V2G and pays the highest electricity prices.

Overall, these results demonstrate that the SMPC model strikes a balance between service quality, operational efficiency and energy costs. It delivers shorter wait times than deterministic or reactive baselines, avoids the conservatism of robust optimization and captures meaningful grid revenues by exploiting dynamic electricity prices. While the improvements over deterministic and reactive baselines are consistent across all metrics, the absolute performance gaps remain moderate. This is partly explained by the relatively short 10-minute control interval used in the SMPC framework. Frequent re-optimization introduces strong closed-loop feedback, allowing the system to correct forecast errors rapidly and thereby reducing the marginal benefit of long-horizon anticipation.

This highlights an important structural property of high-frequency receding horizon control: as the update frequency increases, feedback dominates forecasting accuracy, and the incremental value of multi-stage stochastic optimization decreases. In contrast, for longer control intervals—where decisions must be committed further in advance and corrective adjustments are delayed—the relative advantage of anticipatory stochastic optimization is expected to increase.

6.4.1. SMPC component analysis

To isolate the contribution of each component in the proposed framework, we conducted a component analysis study of the SMPC with four configurations: (1) the full SMPC baseline; (2) SMPC without chance constraints, where energy consumption and travel time are set to their expected values; (3) SMPC without the scenario tree, using a single expected-value path; and (4) SMPC with a reduced robust horizon of 2 stages. [Table 7](#) reports operational performance for San Francisco under the peak pricing scenario.

The scenario tree is the most impactful component: removing it increases the median wait time from 229s to 240s and the 95th percentile from 548s to 571s, closely matching the DMP baseline, which confirms that branching over demand scenarios is the key differentiating element of the framework. Removing chance constraints leaves the median unchanged but degrades the 95th percentile (548→568s), indicating that chance constraints primarily improve tail reliability by preventing dispatch of vehicles with insufficient charge. Reducing the robust horizon to 2 stages produces a moderate but consistent degradation across both metrics (median 229→236s, 95% 548→561s), demonstrating graceful degradation with a shorter planning horizon while retaining most of the performance gain. Together, these results confirm that each component contributes distinctly: the scenario tree drives average performance, chance constraints improve service reliability, and the robust horizon balances solution quality with computational tractability.

Table 7
SMPC component analysis (San Francisco, peak pricing).

Configuration	Method	Median [s]	95% [s]	Total Dist./Veh. [km]	Reb. Dist./Veh. [km]
San Francisco	Full SMPC	229	548	414.1	108.5
	No chance constraints	229	568	410.6	103.3
	No scenario tree (single path)	240	571	408.3	101.7
	Robust horizon = 2	236	561	410.3	102.5

Table 8
Operational performance in Chicago (fleet size 1600) and San Francisco (fleet size 500) under four control strategies. Metrics include mean, median (50% quantile), and 95% percentile wait times (WT), as well as total and rebalance distance per vehicle.

City	Method	Median [s]	95% [s]	Total Dist./Veh. [km]	Reb. Dist./Veh. [km]
Chicago	RB	272	955	456.1	4.6
	DMP	215	872	490.0	65.1
	RCC	220	822	510.1	96.6
	SMPC	214	767	500.7	70.1
San Francisco	RB	360	670	401.1	3.8
	DMP	240	584	407.1	99.5
	RCC	254	559	431.6	124.8
	SMPC	229	548	414.1	108.5

Table 9
Average electricity price paid per kilowatt-hour (euro/kWh) and corresponding total fleet vehicle-to-grid (V2G) revenue (euro) under flat, peak, and solar pricing schemes for Chicago and San Francisco.

City	Method	Flat Price [euro/kWh]	Flat V2G Rev. [euro]	Peak Price [euro/kWh]	Peak V2G Rev. [euro]	Solar Price [euro/kWh]	Solar V2G Rev. [euro]
Chicago	RB	0.06	0.0	0.083	0.00	0.09	0.00
	DMP	0.06	0.0	0.078	565	0.07	981
	RCC	0.06	0.0	0.076	547	0.08	944
	SMPC	0.06	0.0	0.075	563	0.07	975
San Francisco	RB	0.06	0.0	0.08	0.00	0.090	0.00
	DMP	0.06	0.0	0.07	290	0.070	515
	RCC	0.06	0.0	0.08	285	0.074	453
	SMPC	0.06	0.0	0.069	291	0.068	509

Table 10
Overview of the experimental setup for the mobility-on-demand simulations in San Francisco and Chicago. The different vehicle types corresponds to different energy efficiency (low, medium and high).

Factor	San Francisco	Chicago
Vehicle Type	1 (low), 2 (medium), 3 (high)	1 (low), 2 (medium), 3 (high)
Battery Size (kWh)	40, 70, 100	40, 70, 100
Fleet Size (vehicles)	450, 500, 550	1400, 1600, 1800
Number of Chargers	100, 150, 200	200, 300, 400

6.5. Experimental setup for sensitivity analysis

This section investigates the impact of battery size and vehicle energy efficiency on the performance of urban EAMoD systems in San Francisco and Chicago. We analyze three distinct vehicle types, representing high, medium, and low energy efficiency, and consider three battery capacities: 40, 70, and 100 kWh (Table 10). Vehicle parameters, including those determining efficiency, are summarized in Table 11, based on the energy consumption model from Basso et al. (2021). Vehicle weights listed in the table exclude battery mass, which is added separately assuming a specific energy of 268 Wh/kg Günter and Wassiliadis (2022). Fig. 8 illustrates the energy consumption rates of the three vehicle types with a 40 kWh battery, excluding auxiliary loads.

Table 11

Vehicle parameters used to model energy consumption. Symbols, descriptions, and units are listed in the first three columns.

Symbol	Description	Unit	Vehicle 1	Vehicle 2	Vehicle 3
m	Weight	kg	1500	1700	1900
η^+	Efficiency (acceleration)	-	0.90	0.85	0.80
η^-	Efficiency (deceleration)	-	0.90	0.85	0.80
η^{nomial}	Efficiency (constant speed)	-	0.90	0.85	0.80
c_r	Rolling resistance coefficient	-	0.0072	0.009	0.0108
c_d	Drag coefficient	-	0.22	0.25	0.30
p_{aux}	Auxiliary power	W	800	900	1000
p_{auxidle}	Auxiliary idle power	W	100	100	100
A	Frontal area	m ²	2.2	2.5	2.8




Fig. 8. Energy consumption of the different vehicles without auxiliary power and with a 40 kWh battery.

6.6. Sensitivity analysis of battery size and energy efficiency (San Francisco)

Battery capacity significantly influences system performance, though the extent depends on vehicle efficiency (Table 12). For Vehicle 1, the most energy-efficient option, wait times remain nearly constant across all battery sizes. Increasing the battery from 40 kWh (median 220s) to 100 kWh (median 219s) yields virtually no change in performance. Similarly, for vehicle 1 the 75th and 95th percentile wait times remain below 450s across all capacities. This indicates that Vehicle 1's service quality is not constrained by battery size, even with the smallest pack.

Vehicles 2 and 3 are more sensitive to battery size. For Vehicle 2, the median wait time decreases slightly from 225 s at 40 kWh to 217 s at 70 kWh, while the share of passengers waiting more than 600 s declines from 154 to 151. Increasing the battery to 100 kWh provides only marginal improvements, with the median wait time rising slightly to 219 s and the number of passengers waiting more than 600 s falling to 144. This suggests that a 70 kWh battery is sufficient to mitigate charging as a performance bottleneck for Vehicle 2, with further increases offering diminishing returns. For Vehicle 3, the least efficient option, the effect of battery size is far more pronounced. With 40 kWh, the system exhibits very poor performance, with a median wait time of 325s, a 95th percentile of 2625s, and more than 3600 passengers experiencing delays over 600s. Upgrading to 70 kWh brings dramatic improvements: the median drops to 219s, the 95th percentile to 432s, and delays over 600s fall by more than 95% to just 145. Increasing to 100 kWh does not significantly improve service further, with a median of 222s and 149 delays over 600s. This highlights that for inefficient vehicles, battery size is a critical driver of feasibility, with performance plateauing around 70 kWh.


These results yield two key insights. First, energy-efficient vehicles like Vehicle 1 exhibit lower sensitivity to battery size, enabling effective performance with smaller, lighter battery packs. This directly contributes to overall system efficiency and potentially lower operational costs. Second, while less efficient vehicles require larger batteries to achieve acceptable service quality, the performance gains tend to plateau around 70 kWh. Importantly, these service dynamics also shape revenues from V2G services. Efficient vehicles consistently generate higher V2G revenues, since their lower consumption leaves more capacity for grid interaction. For example, Vehicle 1 earns 304 euro per vehicle at 40 kWh and rises to 390 euro at 100 kWh, the highest across all types. Vehicle 3, by contrast, produces the lowest revenues (e.g., 260 euro at 40 kWh, 360 euro at 100 kWh), as its higher driving demand limits grid availability. Larger batteries generally boost revenues across all vehicles, but the marginal gains diminish at the upper end: Vehicle 2 increases from 285 euro at 40 kWh to 342 euro at 70 kWh, while the additional jump to 386 euro at 100 kWh is relatively small. Thus, the same diminishing returns seen in wait times and charging behavior are mirrored in V2G participation.

These findings are consistently reflected in other system-level metrics. Fig. 9 demonstrate that increasing fleet size and charger availability effectively reduce wait times, particularly for the less efficient Vehicle 3 and some for Vehicle 2. Vehicle 1, however,


Table 12

Summary of wait time and per-vehicle distance by battery capacity and vehicle type for San Francisco. The “Wait > 600s” column reports the number of trips with wait times exceeding 600 s and the corresponding percentage of total trip requests ($n = 11,453$).

Vehicle	SoC [kWh]	Median [s]	75% [s]	95% [s]	Wait > 600s [count (%)]	TotDist/Veh [km]	RebalDist/Veh [km]	V2G Total [€]
1	40	220	301	440	137 (1.2%)	240	79	304
2	40	225	309	452	154 (1.3%)	206	42	285
3	40	325	1129	2625	3643 (31.8%)	216	37	260
1	70	219	296	432	145 (1.3%)	284	123	347
2	70	217	296	434	151 (1.3%)	241	80	342
3	70	219	299	432	145 (1.3%)	209	47	316
1	100	219	303	443	134 (1.2%)	291	129	390
2	100	219	303	447	144 (1.3%)	292	130	386
3	100	222	306	451	149 (1.3%)	190	32	360



(a) Median (circle) and 75th percentile (triangle) of wait times across fleet size



(b) Median (circle) and 75th percentile (triangle) of wait times across number of chargers

Fig. 9. Comparison of wait time metrics in San Francisco. Fig. A shows metrics across different fleet sizes and vehicle types with 100 chargers. Fig. B shows metrics across different number of charger and vehicle types with a fleet size of 450 vehicles. All plots assume a 40 kWh battery capacity. For clarity, 75th percentile wait times for Vehicle 3 are not shown in the plots, but their values are as follows: for varying *fleet size* (450, 500, 550 vehicles), the Q75 values are **1129 s**, **610 s**, and **590 s**, respectively; for varying *number of chargers* (100, 150, 200), the Q75 values are **1129 s**, **654 s**, and **650 s**, respectively.

continues to exhibit stable performance, characterized by lower 75th percentile and minimal sensitivity to infrastructure constraints. Furthermore, Vehicle 3 has the highest energy consumption and consequently the greatest charging demand, especially during peak hours, see Fig. 10a. Conversely, Vehicle 1 maintains a smoother, more aligned charging-consumption profile throughout the day.

Fig. 11a reports the estimated number of daily charge cycles (defined as charging from 20% to 80% state of charge) across vehicle types and battery sizes. As anticipated, larger batteries consistently require fewer charging cycles. Vehicle 1 consistently exhibits the lowest and most stable charging frequency, with a median of just under one cycle per day even with a 40 kWh battery. In contrast, Vehicles 2 and 3 demonstrate higher and more variable charging frequencies, particularly with smaller battery configurations. For instance, Vehicle 3 with a 40 kWh pack exhibits a median of just under 1.5 cycles per day and substantial variability in its charging behavior. Increasing the battery capacity to 70 kWh or 100 kWh nearly halves the median number of cycles and significantly narrows the distribution of charging events.

These observed differences in charging frequency directly correlate with wait time performance. Frequent charging, especially for Vehicle 3 with small batteries, leads to service delays as vehicles become temporarily unavailable for passenger service. Larger batteries effectively mitigate both the number of charge cycles and the occurrence of such extended delays.

In summary, this analysis highlights that efficient vehicles can deliver consistent service with smaller batteries, while less efficient vehicles necessitate larger battery packs to avoid degraded performance. Reducing charge cycle frequency is paramount not only for prolonging battery lifespan but also for ensuring stable and responsive EAMoD service. These results underscore that optimal vehicle selection for urban EAMoD systems must consider not only immediate performance metrics but also long-term operational sustainability. Ultimately, if a vehicle is sufficiently energy efficient, a 40 kWh battery offers a compelling trade-off between cost, performance, and battery longevity.




Fig. 10. Comparison of hourly energy charged and consumed across three different vehicle types with 40 kWh battery capacity for (a) San Francisco and 70 kWh for (b) Chicago. Each line represents the total energy (in kWh) aggregated across all vehicles in a fleet. Charging events are shown with dashed lines, and consumption with solid lines.




Fig. 11. Observed charge cycles (20–80% depth of discharge) for different vehicle types (1, 2, 3) and battery sizes (40, 70, 100 kWh) in (a) San Francisco and (b) Chicago. These distributions highlight the combined influence of vehicle characteristics and geographical location on battery charging patterns.

6.7. Sensitivity analysis of battery size and energy efficiency (Chicago)

We now examine how battery size and vehicle energy efficiency affect system performance in Chicago, which presents a different urban context from San Francisco. The same three vehicle types and battery capacities (40, 70, and 100 kWh) are considered, except for Vehicle 3 with a 40 kWh battery, which is excluded because its limited capacity cannot sustain viable operation in Chicago.

Energy demand in Chicago is more concentrated during peak periods than in San Francisco, as shown in Fig. 10b. Less efficient vehicles (Vehicle 3 in particular) exhibit a pronounced mismatch between energy consumption and charging, leading to steep late-day charging requirements. Vehicle 1 maintains the most balanced profile, with low and stable energy use throughout the day. This stability also allows Vehicle 1 to contribute more consistently to V2G services, especially at larger battery sizes.

Battery size has a substantial influence on system performance in Chicago, with less efficient vehicles showing the greatest sensitivity to energy storage capacity. As reported in Table 13, Vehicle 1 delivers consistently low wait times across all battery sizes, even under Chicago's higher demand. The median wait time improves modestly from 214 s at 40 kWh to 203 s at both 70 and 100 kWh. The effect is more visible in the upper percentiles: the 75th percentile drops from 329 s to 258 s between 40 and 70 kWh and remains unchanged at 100 kWh, indicating that a moderate battery increase is sufficient to reduce occasional peak delays. The 95th percentile falls from 1302 s at 40 kWh to 673 s at both 70 and 100 kWh, cutting the longest delays by nearly half. The number of trips with

Table 13

Summary of wait time and per-vehicle distance by battery capacity and vehicle type for Chicago. The “Wait > 600s” column reports the number of trips with wait times exceeding 600 s and the corresponding percentage of total trip requests ($n = 32,406$).

Vehicle	SoC [kWh]	Median [s]	75% [s]	95% [s]	Wait > 600s [count (%)]	TotDist/Veh [km]	RebalDist/Veh [km]	V2G Total [€]
1	40	214	329	1302	4392 (13.6%)	512.3	67.1	572
2	40	293	2129	9772	12,176 (37.6%)	461.7	52.1	563
1	70	203	258	673	1911 (5.9%)	500.7	70.1	563
2	70	221	426	2919	6606 (20.4%)	510.0	60.6	563
3	70	316	5652	18,175	13,113 (40.5%)	432.6	50.9	559
1	100	203	258	673	1941 (6.0%)	492.0	63.5	580
2	100	204	266	692	2004 (6.2%)	504.2	68.8	562
3	100	237	1072	10,830	9912 (30.6%)	466.4	52.8	550




Fig. 12. Comparison of wait time metrics in Chicago. Fig. a shows metrics across different fleet sizes and vehicle types with 200 chargers. Fig. b shows metrics across different number of charger and vehicle types with a fleet size of 1400 vehicles. All plots assume a 70 kWh battery capacity. For clarity, 75th percentile wait times for Vehicle 3 are not shown in the plots, but their values are as follows: for varying *fleet size* (1400, 1600, 1800 vehicles), the Q75 values are **5651 s**, **2180 s**, and **969 s**, respectively; for varying *number of chargers* (200, 300, 400), the Q75 values are **5651 s**, **1172 s**, and **932 s**, respectively.

wait times exceeding 600 s is likewise more than halved, dropping from 4392 to approximately 1900. These results show that while Vehicle 1 already performs well at the median, larger batteries primarily improve service reliability by reducing infrequent but severe delays. Vehicle 1 also achieves the highest V2G revenues, ranging from €572 at 40 kWh to €580 at 100 kWh, confirming that efficiency and moderate battery sizing jointly support both passenger service and grid interaction.

Vehicle 2, by contrast, depends much more strongly on battery size. Its median wait time decreases from 293 s at 40 kWh to 221 s at 70 kWh and 204 s at 100 kWh. The tail behaviour is even more striking: the 75th percentile shrinks from 2129 s at 40 kWh to 426 s at 70 kWh and 266 s at 100 kWh, while the 95th percentile drops from 9772 s to 2919 s and then to 692 s. This represents a shift from highly unstable to near-stable performance. The number of trips exceeding 600 s falls from 12 176 at 40 kWh to 6606 at 70 kWh and 2004 at 100 kWh. Although more efficient than Vehicle 3, Vehicle 2 remains sensitive to energy availability and benefits considerably from battery upgrades. Its V2G revenues, however, stay nearly flat across capacities (€563 at 40–70 kWh, €562 at 100 kWh), indicating that most of its stored energy is consumed by travel demand rather than made available for grid services.

Vehicle 3, the least efficient type, is the most sensitive to battery capacity and consistently struggles under Chicago’s high demand. At 70 kWh, its median wait time is 316 s, but the upper tail reveals severe congestion: the 75th percentile reaches 5652 s and the 95th percentile exceeds 18 000 s, with over 13 000 trips delayed by more than 10 minutes. A 100 kWh battery offers partial relief, reducing the median to 237 s, the 75th percentile to 1072 s, and the 95th percentile to 10 830 s. Long-wait instances drop to 9912, but this remains far too high for acceptable service. These figures illustrate the fundamental limitations of low-efficiency vehicles in dense urban settings: even substantial battery investments cannot close the performance gap with more efficient alternatives. Vehicle 3 also generates the lowest V2G revenues (€559 at 70 kWh and €550 at 100 kWh), as its high consumption leaves little surplus energy for grid participation.

Fig. 12 shows how infrastructure scale affects performance. As fleet size increases from 1400 to 1800, all vehicles benefit, but the reduction is most pronounced for Vehicles 2 and 3. Vehicle 3, for instance, sees its median wait time fall from over 300 s to around

240 s. Similarly, increasing the number of chargers from 200 to 400 lowers both median wait times and the 75th percentile, with the largest gains again accruing to less efficient vehicles.

Charging frequency data (Fig. 11b) follow similar trends as in San Francisco. Vehicle 1 requires just over one cycle per day even with a 40 kWh battery, whereas Vehicles 2 and 3 charge more frequently and with greater variability, especially at smaller battery sizes. Vehicle 2 at 40 kWh has a median of about 1.2 cycles per day, dropping to around 0.7 at 70 kWh. Vehicle 3 consistently records the highest cycle counts, reflecting its greater energy demand. Frequent charging not only limits vehicle availability for passenger trips but also reduces the window for stable V2G participation.

The Chicago results reinforce the San Francisco findings while exposing sharper contrasts due to higher demand intensity. Vehicle 1 remains robust across battery sizes and system configurations, whereas Vehicles 2 and 3 require both larger batteries and more infrastructure to maintain acceptable service levels. Pairing efficient vehicles with moderate battery capacities yields strong, scalable performance; less efficient platforms demand combined investment in capacity and infrastructure to avoid service degradation. From a V2G perspective, Vehicle 1 is again the most attractive option, combining stable service with the highest revenues. Vehicles 2 and 3 offer limited V2G potential because their stored energy is largely consumed by travel demand, confirming that vehicle efficiency is the primary determinant of grid-service opportunities in high-demand urban environments.

7. Conclusion

This paper presented a combined stochastic and robust model predictive control framework for electric autonomous mobility-on-demand (EAMoD) systems under multiple sources of uncertainty. By integrating Bayesian Neural Field based probabilistic forecasting with multi-stage stochastic optimization and robust chance constraints, the proposed framework jointly optimizes dispatch, rebalancing, charging, and V2G decisions in a unified formulation.

A key methodological contribution is the tailored Nested Benders Decomposition algorithm, which exploits the scenario tree structure to solve large-scale multi-stage mixed-integer problems with temporally correlated uncertainties. This enables real-time applicability despite the exponential growth of the extensive-form formulation. Unlike standard SAA formulations that typically assume stagewise independence, the proposed approach preserves temporal dependence across demand and infrastructure uncertainties, which is essential in EAMoD systems.

Simulation results on San Francisco and Chicago demonstrate consistent improvements over deterministic, reactive, and robust baselines, including up to 36% reduced passenger waiting times, 27% lower rebalancing distance, and close to 35% electricity cost savings through V2G participation. The Bayesian Neural Field forecasting model was shown to outperform established baselines in both point accuracy and calibrated uncertainty quantification, providing the probabilistic inputs essential for effective scenario tree construction. Importantly, the analysis reveals a structural insight: the relative value of multi-stage stochastic anticipation depends on the control update frequency. When receding horizon updates are frequent, closed-loop feedback mitigates forecast errors, reducing but not eliminating the marginal benefit of anticipatory optimization. For longer decision intervals or stronger temporal coupling in uncertainties, the advantage of stochastic control is expected to grow.

Overall, the proposed framework provides a scalable and computationally tractable pathway for integrating predictive uncertainty modeling and energy-aware fleet control in next-generation autonomous mobility systems. Future research may explore scenario tree construction based on the nested distance (Pflug, 2010), which provides formal approximation guarantees by accounting for the filtration structure of the stochastic process. Additionally, the current price-based adaptive terminal cost is a tractable heuristic that does not account for demand forecasts, fleet state, or spatial factors; designing a more principled terminal value remains a key open challenge. Promising directions include approximate dynamic programming or reinforcement learning to learn state-dependent terminal value functions from offline simulation (Bertsekas, 2024; Rosolia and Borrelli, 2017; Baltussen et al., 2025), and hierarchical or bi-level optimization architectures that couple a coarse long-horizon planner with the short-horizon SMPC (Amini et al., 2019). Multi-scale time discretizations and alternative scenario sampling methods also warrant further investigation to improve long-horizon foresight within the receding horizon framework. Finally, the framework could be extended to endogenous pricing, network congestion effects, or co-optimization of charging infrastructure and fleet composition.

CRedit authorship contribution statement

Sten Elling Tingstad Jacobsen: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization; **Balázs Kulcsár:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Conceptualization; **Anders Lindman:** Writing – review & editing, Software, Resources, Project administration, Funding acquisition, Conceptualization.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used ChatGPT and DeepL exclusively to improve the readability and language of the manuscript. The scientific content, ideas, and conclusions are entirely the authors' own. After using these tools, the authors carefully reviewed and edited the text, and they take full responsibility for all content of the published article.

Acknowledgements

This work was co-funded by [Vinnova](#), Sweden through the project: Simulation, analysis and modeling of future efficient traffic systems. This work was in part supported by the Transport Area of Advance within Chalmers University of Technology. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS) at Chalmers e-Commons partially funded by the Swedish Research Council through grant agreement no 2022-06725.

Data availability

Data will be made available on request.

References

- Ameer, H., Wang, Y., Fan, X., Chen, Z., 2025. Hybrid optimization of EV charging station placement and pricing using bender's decomposition and NSGA-II algorithm. *Appl. Energy* 397, 126385. <https://doi.org/10.1016/j.apenergy.2025.126385>
- Amini, M.R., Kolmanovsky, I., Sun, J., 2019. Robust hierarchical MPC for handling long horizon demand forecast uncertainty with application to automotive thermal management. In: 2019 IEEE 58th Conference on Decision and Control (CDC). IEEE, pp. 6694–6699.
- Ansariipoor, A.H., Oliveira, F.S., Liret, A., 2016. Recursive expected conditional value at risk in the fleet renewal problem with alternative fuel vehicles. *Transport. Res. Part C* 65, 156–171.
- Baltussen, T., Orrico, C.A., Katriniok, A., Heemels, W., Krishnamoorthy, D., 2025. Value function approximation for nonlinear MPC: learning a terminal cost function with a descent property. In: 2025 IEEE 64th Conference on Decision and Control (CDC). IEEE, pp. 5507–5514.
- Bao, Y., Chan, K.J., Mesbah, A., Velni, J.M., 2023. Learning-based adaptive-scenario-tree model predictive control with improved probabilistic safety using robust bayesian neural networks. *Int. J. Robust Nonlinear Control* 33 (5), 3312–3333.
- Basso, R., Kulcsár, B., Sanchez-Diaz, I., 2021. Electric vehicle routing problem with machine learning for energy prediction. *Transp. Res. Part B* 145, 24–55.
- Batty, M., Axhausen, K.W., Giannotti, F., Pozdnoukhov, A., Bazzani, A., Wachowicz, M., Ouzounis, G., Portugali, Y., 2012. Smart cities of the future. *Eur. Phys. J. Spec. Topics* 214, 481–518. <https://doi.org/10.1140/epjst/e2012-01703-3>
- Beltrán, F., Finardi, E.C., de Oliveira, W., 2021. Two-stage and multi-stage decompositions for the medium-term hydrothermal scheduling problem: a computational comparison of solution techniques. *Int. J. Electr. Power Energy Syst.* 127, 106659.
- Benders, J.F., 1962. Partitioning procedures for solving mixed-variables programming problems. *Numerische Math.* 4 (1), 238–252.
- Bertsekas, D.P., 2024. Model predictive control and reinforcement learning: a unified framework based on dynamic programming. *IFAC-PapersOnLine* 58 (18), 363–383. <https://doi.org/10.1016/j.ifacol.2024.09.056>
- Birge, J.R., 1985. Decomposition and partitioning methods for multistage stochastic linear programs. *Operat. Res.* 33 (5), 989–1007.
- Birge, J.R., Louveaux, F., 2011. *Introduction to Stochastic Programming*. Springer New York.
- Chen, R., Luedtke, J., 2022. On generating lagrangian cuts for two-stage stochastic integer programs. *INFORMS J. Comput.* 34 (4), 2332–2349.
- Chen, Z.-L., Powell, W.B., 1999. Convergent cutting-plane and partial-sampling algorithm for multistage stochastic linear programs with recourse. *J. Optim. Theor. Appl.* 102 (3), 497–524.
- City of Chicago, 2023. Taxi trips (2013 - 2023). https://data.cityofchicago.org/Transportation/Taxi-Trips-2013-2023-/wrvz-psew/about_data. Accessed: 2024-10-01.
- Dempster, M. A.H., Thompson, R.T., 1998. Parallelization and aggregation of nested benders decomposition. *Ann. Oper. Res.* 81 (0), 163–188.
- Donohue, C.J., Birge, J.R., 2006. The abridged nested decomposition method for multistage stochastic linear programs with relatively complete recourse. *Algor. Oper. Res.* 1 (1), 20–30.
- Dupačová, J., Gröwe-Kuska, N., Römis, W., 2003. Scenario reduction in stochastic programming. *Math. Program.* 95, 493–511.
- Ermoliev, Y.M., 1983. Stochastic quasi-gradient methods and their application to systems optimization. *Stochastics* 9 (1–2), 1–36. <https://doi.org/10.1080/17442508308833246>
- Faridimehr, S., Venkatachalam, S., Chinnam, R.B., 2018. A stochastic programming approach for electric vehicle charging network design. *IEEE Trans. Intell. Transp. Syst.* 20 (5), 1870–1882.
- Chee, C.H., Neo, H., 2018. Five big challenges facing big cities of the future. *World Econ. Forum* <https://www.weforum.org>.
- Gade, D., Küçükyavuz, S., Sen, S., 2014. Decomposition algorithms with parametric gomory cuts for two-stage stochastic integer programs. *Math. Program.* 144 (1), 39–64.
- Gao, J., Li, S., 2024. Charging autonomous electric vehicle fleet for mobility-on-demand services: plug in or swap out? *Transport. Res. Part C* 158, 104457.
- Günter, F.J., Wassiliadis, N., 2022. State of the art of lithium-ion pouch cells in automotive applications: cell teardown and characterization. *J. Electrochem. Soc.* 169 (3), 030515.
- Gurobi Optimization, L., 2022. *Gurobi Optimizer Reference Manual*. <https://www.gurobi.com>.
- He, S., Zhang, Z., Han, S., Pepin, L., Wang, G., Zhang, D., Stankovic, J.A., Miao, F., 2023. Data-driven distributionally robust electric vehicle balancing for autonomous mobility-on-demand systems under demand and supply uncertainties. *IEEE Trans. Intell. Transp. Syst.* 24 (5), 5199–5215.
- Higle, J.L., Sen, S., 1991. Stochastic decomposition: an algorithm for two-stage linear programs with recourse. *Math. Oper. Res.* 16 (3), 650–669.
- Hochreiter, R., Pflug, G.C., 2007. Financial scenario generation for stochastic multi-stage decision processes as facility location problems. *Ann. Oper. Res.* 152, 257–272.
- Hoffman, A.J., Kruskal, J.B., 2010. *Integral Boundary Points of Convex Polyhedra. 50 Years of Integer Programming 1958–2008: From the Early Years to the State-of-the-Art*. Springer Berlin Heidelberg, Berlin, Heidelberg. pp. 49–76.
- Horejšová, M., Vitali, S., Kopa, M., Moriggia, V., 2020. Evaluation of scenario reduction algorithms with nested distance. *Comput. Manag. Sci.* 17 (2), 241–275.
- Horni, A., Nagel, K., Axhausen, K.W., 2016. *The Multi-Agent Transport Simulation MATSim*. Ubiquity Press, London, GBR.
- Høyland, K., Wallace, S.W., 2001. Generating scenario trees for multistage decision problems. *Manag. Sci.* 47 (2), 295–307.
- Iglesias, R., Rossi, F., Wang, K., Hallac, D., Leskovec, J., Pavone, M., 2018. Data-driven model predictive control of autonomous mobility-on-demand systems. In: 2018 IEEE International Conference on Robotics and Automation (ICRA). IEEE, pp. 6019–6025.
- Iglesias, R., Rossi, F., Zhang, R., Pavone, M., 2019. A BCMP network approach to modeling and controlling autonomous mobility-on-demand systems. *Int. J. Robot. Res.* 38 (2–3), 357–374.
- Infanger, G., Morton, D.P., 1996. Cut sharing for multistage stochastic linear programs with interstage dependency. *Math. Program.* 75 (2), 241–256.
- Jacobsen, S. E.T., Lindman, A., Kulcsár, B., 2023. A predictive chance constraint rebalancing approach to mobility-on-demand services. *Commun. Transp. Res.* 3, 100097.
- Jin, X., Han, J., 2017. *K-Means Clustering*. Springer US, Boston, MA. pp. 695–697. https://doi.org/10.1007/978-1-4899-7687-1_431
- Kchaou-Boujelben, M., 2021. Charging station location problem: a comprehensive review on models and solution approaches. *Transport. Res. Part C* 132, 103376.
- Kovacevic, R.M., Pichler, A., 2015. Tree approximation for discrete time stochastic processes: a process distance approach. *Ann. Oper. Res.* 235 (1), 395–421.
- Lei, T., Guo, S., Qian, X., Gong, L., 2022. Understanding charging dynamics of fully-electrified taxi services using large-scale trajectory data. *Transport. Res. Part C* 143, 103822.
- Lucia, S., Finkler, T., Engell, S., 2013. Multi-stage nonlinear model predictive control applied to a semi-batch polymerization reactor under uncertainty. *J. Process Control* 23 (9), 1306–1319.

- Miao, F., He, S., Pepin, L., Han, S., Hendawi, A., Khalefa, M.E., Stankovic, J.A., Pappas, G., 2021. Data-driven distributionally robust optimization for vehicle balancing of mobility-on-demand systems. *ACM Trans. Cyber-Phys. Syst.* 5 (2), 1–27.
- Murphy, J., 2013. Benders, nested benders and stochastic programming: an intuitive introduction. 1312.3158arXiv preprint arXiv:1312.3158.
- Nature Editorial, Want a sustainable future? Then look to the world's cities. *Nature*, 620 (7975):697, August 2023. Editorial. <https://doi.org/10.1038/d41586-023-02597-3>
- Nieto-Isaza, S., Fontaine, P., Minner, S., 2022. The value of stochastic crowd resources and strategic location of mini-depots for last-mile delivery: a benders decomposition approach. *Transp. Res. Part B* 157, 62–79. <https://doi.org/10.1016/j.trb.2021.12.014>
- Nord Pool, 2025. Day-ahead market prices. <https://data.nordpoolgroup.com/auction/day-ahead/prices?deliveryDate=2025-01-25¤cy=SEK&aggregation=DeliveryPeriod&deliveryAreas=SE2>. Accessed: 2025-02-03.
- Paparella, F., Hofman, T., Salazar, M., 2024. Electric autonomous mobility-on-demand: jointly optimal vehicle design and fleet operation. *IEEE Trans. Intell. Transp. Syst.* 25 (11), 17054–17065. <https://doi.org/10.1109/ITITS.2024.3428569>
- Parra, G., Tobar, F., 2017. Spectral mixture kernels for multi-output gaussian processes. *Adv. Neural Inform. Process. Syst.* 30.
- Pereira, M. V.F., Pinto, L.M., 1991. Multi-stage stochastic optimization applied to energy planning. *Math. Program.* 52 (1), 359–375.
- Pflug, G.C., 2010. Version-independence and nested distributions in multistage stochastic optimization. *SIAM J. Optim.* 20 (3), 1406–1420.
- Piorkowski, M., Sarafijanovic-Djukic, N., Grossglauser, M., 2022. Cawdad epfl/mobility. Dataset. <https://doi.org/10.15783/C7J010>
- Qin, Z.T., Zhu, H., Ye, J., 2022. Reinforcement learning for ridesharing: an extended survey. *Transport. Res. Part C* 144, 103852.
- Rawlings, J.B., Mayne, D.Q., Diehl, M., 2017. *Model Predictive Control: Theory, Computation, and Design*.
- Richards, A., How, J., 2005. Mixed-integer programming for control. In: *Proceedings of the 2005, American Control Conference, 2005. IEEE*, pp. 2676–2683.
- Robbins, H., Monro, S., 1951. A stochastic approximation method. *Annals Math. Stat.* 22 (3), 400–407.
- Rosolia, U., Borrelli, F., 2017. Learning model predictive control for iterative tasks. a data-driven control framework. *IEEE Trans. Autom. Control* 63 (7), 1883–1896.
- Ruch, C., Gächter, J., Hakenberg, J., Frazzoli, E., 2020. The + 1 method: model-free adaptive repositioning policies for robotic multi-agent systems. *IEEE Trans. Network Sci. Eng.* 7 (4), 3171–3184.
- Ruch, C., Hör, S., Frazzoli, E., 2018. Amodeus, a simulation-based testbed for autonomous mobility-on-demand systems. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 3639–3644. <https://doi.org/10.1109/ITSC.2018.8569961>
- Ruszczynski, A.P., 1986. A regularized decomposition method for minimizing a sum of polyhedral functions. *Math. Program.* 35 (3), 309–333.
- Saad, F., Burnim, J., Carroll, C., Patton, B., Köster, U., A. Saurous, R., Hoffman, M., 2024. Scalable spatiotemporal prediction with bayesian neural fields. *Nature Commun.* 15 (1), 7942.
- Sen, S., Hige, J.L., 1999. An introductory tutorial on stochastic linear programming models. *Interfaces* 29 (2), 33–61.
- Sen, S., Zhou, Z., 2014. Multistage stochastic decomposition: a bridge between stochastic programming and approximate dynamic programming. *SIAM J. Optim.* 24 (1), 127–153. <https://doi.org/10.1137/120864854>
- Shapiro, A., Dentcheva, D., Ruszczyński, A., 2011. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM. 2 edition.
- Spieser, K., Treleaven, K., Zhang, R., Frazzoli, E., Morton, D., Pavone, M., 2014. Toward a systematic approach to the design and evaluation of automated mobility-on-demand systems: a case study in singapore. In: *Road Vehicle Automation*. Springer, pp. 229–245.
- Tsao, M., Iglesias, R., Pavone, M., 2018. Stochastic model predictive control for autonomous mobility on demand. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, pp. 3941–3948.
- Turan, B., Pedarsani, R., Alizadeh, M., 2020. Dynamic pricing and fleet management for electric autonomous mobility on demand systems. *Transport. Res. Part C* 121, 102829.
- United Nations Development Programme, 2017. Rapid urbanisation: opportunities and challenges to improve the well-being of societies. <https://hdr.undp.org>.
- United Nations Development Programme (UNDP), 2020. Undp issues brief on urban climate resilience. UNDP Issues Brief. Published November 2, 2020. <https://www.undp.org/publications/undp-issues-brief-urban-climate-resilience>.
- van Ackooij, W., Warin, X., 2020. On conditional cuts for stochastic dual dynamic programming. *EURO J. Comput. Optim.* 8 (2), 173–199. <https://doi.org/10.1007/s13675-020-00123-y>
- Van Slyke, R.M., Wets, R. J.B., 1969. L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM J. Appl. Math.* 17 (4), 638–663.
- Wassiliadis, N., Schneider, J., Frank, A., Wildfeuer, L., Lin, X., Jossen, A., Lienkamp, M., 2021. Review of fast charging strategies for lithium-ion battery systems and their applicability for battery electric vehicles. *J. Energy Stor.* 44, 103306. <https://doi.org/10.1016/j.est.2021.103306>
- Yang, J., Dong, J., Hu, L., 2017. A data-driven optimization-based approach for siting and sizing of electric taxi charging stations. *Transport. Res. Part C* 77, 462–477.
- Yang, J., Levin, M.W., Hu, L., Li, H., Jiang, Y., 2023. Fleet sizing and charging infrastructure design for electric autonomous mobility-on-demand systems with endogenous congestion and limited link space. *Transport. Res. Part C* 152, 104172.
- Zardini, G., Lanzetti, N., Pavone, M., Frazzoli, E., 2022. Analysis and control of autonomous mobility-on-demand systems. *Annu. Rev. Control, Robot., Auton. Syst.* 5, 633–658. <https://doi.org/10.1146/annurev-control-042920-012811>
- Zhang, R., Rossi, F., Pavone, M., 2016. Model predictive control of autonomous mobility-on-demand systems. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 1382–1389.
- Zivot, E., Wang, J., 2006. *Vector Autoregressive Models for Multivariate Time Series*. In: *Modeling Financial Time Series with S-PLUS®*. Springer New York, New York, NY. pp. 369–413.
- Zou, J., Ahmed, S., Sun, X.A., 2019. Stochastic Dual Dynamic Integer Programming. *Math. Program.* 175 (1), 461–502.
- Çalik, H., Fortz, B., 2019. A benders decomposition method for locating stations in a one-way electric car sharing system under demand uncertainty. *Transp. Res. Part B* 125, 121–150. <https://doi.org/10.1016/j.trb.2019.05.004>



Sten Elling Tingstad Jacobsen received the M.S. degree in applied physics from Chalmers University of Technology, Sweden, in 2020, and the Ph.D. degree jointly from Chalmers University of Technology and Volvo Cars AB, Sweden, in 2026. His research interests include optimization, machine learning, control, and their application to transportation systems and the automotive area.



Balázs Kulcsár received the M.Sc. degree in traffic engineering and the Ph.D. degree from the Budapest University of Technology and Economics (BUTE), Budapest, Hungary, in 1999 and 2006, respectively. He has been a Researcher/Post-Doctor with the Department of Control for Transportation and Vehicle Systems, BUTE, the Department of Aerospace Engineering and Mechanics, University of Minnesota, Minneapolis, MN, USA, and the Delft Center for Systems and Control, Delft University of Technology, Delft, the Netherlands. He is currently a Professor with the Department of Electrical Engineering, Chalmers University of Technology, Gothenburg, Sweden. His main research interest focuses on traffic network modeling and optimization.



Anders Lindman received his Ph.D. (2017) in Computational Materials Physics from Chalmers University of Technology (Sweden). He joined Volvo Car Corporation in 2018, where he has been active in the areas of Data & Analytics, Vehicle Energy Efficiency and Future Mobility.