POLITECNICO DI TORINO

Master's Degree in Mechatronic Engineering

Master's Degree Thesis

Digital Twin for Smart Charging Systems

Supervisors

Prof. Giuseppe Carlo CALAFIORE

Dr. Luca AMBROSINO

Candidate

Shakiba SAIF

October 2025

Abstract

The number of electric vehicles (EVs) has increased every year in the recent past, and we expect it to continue growing in the future. EVs lower emissions and reliance on fossil fuels compared to conventional vehicles, but they also increase the strain on the grid of electricity. The fact that many cars frequently charge at the same hours, which can waste opportunities to use, raises costs, and overloads the electricity grid with peaks of energy demand. Nowadays, the most common strategy used to manage the charging of EVs is the so-called First-In First-Served rule (FIFS), which only prioritizes EVs charging based on their arrival time. Smarter algorithms can be instead implemented to make the charging process more efficient, cheaper and faster, depending on our goals. At the same time, we may require techniques that are not just inexpensive but also more environmentally friendly moving towards a greener future.

The main contribution of this thesis is the development of a digital twin platform in Python. Beyond serving as a simulator, our control panel acts as a practical tool for safely testing charging strategies before they are applied in real systems. Thanks to its interactive and easy-to-use design, the platform supports researchers, operators, and policymakers in exploring different approaches and in shaping charging systems that are efficient, sustainable, and responsive to user needs. Technically speaking, we used several libraries such as CVXPY, matplotlib, numpy, pandas, and customtkinter for making our panel complete and professional in order to test any smart charging models desired using different values and scenarios.

Three distinct smart charging models are also presented in this thesis and tested using our simulation platform. The first is the Power Allocation model, which aims to reduce the electricity costs of EVs charging by shifting the power allocation to low-peak periods, without failing the user demand satisfaction. The second is the CO_2 Emission model, which aims to minimize the carbon emission due to the electricity usage, together with a usual cost minimization. The third is the Cohort model, which enables the management of big fleets of EVs by aggregating them into different classes according to their similarities and energy demand. These models collectively address the three primary objectives of smart charging: cost reduction, emission reduction, and efficiency.

Thanks to the easily accessible control panel of our digital twin simulator, we

can virtually apply these smart charging models to evaluate their performances, especially comparing them to the FIFS approach. As expected, the smart charging models show better performances than FIFS both in cost-saving and lowering emissions, depending on the model used. The Cohort model in particular shows how the charging of large groups of vehicles can be managed efficiently in the future.

The control panel allows users to change any parameter like the electricity price used, the fleet size, time resolution, and other power-related constraints parameters. These features allow for an easy comparison of the models by examining how they respond to shifting circumstances. The results are instantly displayed in several plots, and any end-user would be able to analyze them thanks to an AI-generated assistant. Ultimately, our digital twin bridges the gap between theory and deployment, accelerating smarter, greener EV charging. The platform's explainable outputs and what-if analysis help operators and policymakers understanding trade-offs among cost, emissions and efficiency, paving the way for the development of future smart cities.

Acknowledgements

I would like to express my sincere gratitude to Professor Giuseppe C. Calafiore and Dr. Luca Ambrosino for their invaluable guidance, continuous support, and insightful feedback throughout the development of this thesis. Their expertise and encouragement have been fundamental to both my academic and personal growth.
I would also like to express my heartfelt thanks to my husband, whose endless support, patience, and encouragement have been my greatest source of strength during this journey.
My deepest appreciation goes to my family, for their constant love and belief in me, and to all my friends, whose support and kindness have made this experience truly meaningful.
Finally, I sincerely hope that this thesis may serve as a useful reference for future students and researchers, help others in their own academic and professional journeys, and contribute in some small way to making research in this field more impactful and meaningful.

Table of Contents

List of Figures

Chapter 1

Introduction

1.1 Electric Vehicles and Smart Charging

Electric vehicles (EVs), are playing a crucial role in the global shift toward environmentally friendly ways of transportation. Government funding, battery technology advancements, and growing public interest in cleaner mobility are all contributing to their rapidly increasing popularity [1].

Figure 1.1: An electric vehicle charging at a public station [1].

This transformation not only affects how we move but also impacts energy systems and infrastructure.

The evolution of electric vehicles (EVs) over time is shown in Figure 1.2: By early 2025, electric vehicles will account for over 40% of the global market. This demonstrates how electric vehicles are gradually replacing traditional internal combustion engine (ICE) cars that run on gasoline or diesel. A number of factors,

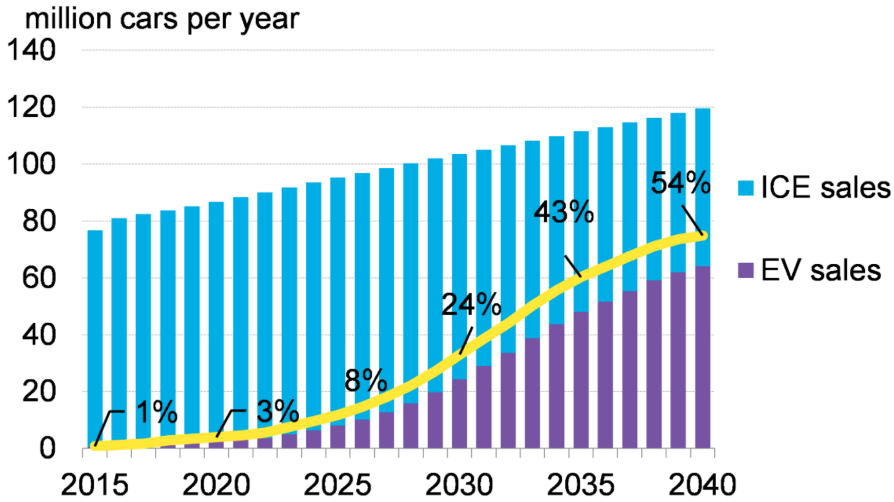

Figure 1.2: Projected growth in global EV market share (2020–2025) [2].

including dropping battery prices and advancements in charging infrastructure that make EVs more feasible for daily use, are contributing to the acceleration.

Rules and policies introduced by governments are equally important. Stricter emission standards, financial rewards for buying EVs, and an increase of charging stations all speed up adoption. At the same time, people are more worried about climate change and want to use less oil and gas. This promotes the use of electric vehicles by both consumers and automakers.

In addition to the rate of adoption, the figure shows the breaking point in the industry, where EVs are moving from a specialized technology to the general public. This change highlights the need to develop efficient charging strategies, such as those discussed in this thesis, to ensure that infrastructure and grid management can satisfy demand.

As EV numbers grow, electricity demand rises especially during peak hours when many users plug in after work. If not managed well, this concentrated demand can overload distribution grids and raise operational costs [3]. Research has shown that independent EV charging could increase peak load by up to 20%, leading to the need for expensive grid upgrades [4]. Simple methods like First In First Served (FIFS) are still commonly used due to their ease of implementation. These strategies ignore key factors, such as time-varying electricity prices, that can lead to ineffective outcomes.

For fixing these issues, we decided to use smart charging techniques. Using real-time measurements such as grid load, energy prices, and renewable output, they move the charging to times when costs are lower or emissions are lower. This ensures electrical stability, increases efficiency, and minimizes the impact on the environment. At the same time, it is important to maintain convenience for users, vehicles must be charged and ready when needed. One of the main reasons behind the optimization models used in this thesis is achieving this balance.

1.2 Literature Review

Optimization Approaches

To make charging EVs better, researchers have been coming to develop a lot of different models. Others try to reduce costs by using state of charge data [5], while others use scheduling strategies that find a balance between user needs and efficiency [6]. Model Predictive Control (MPC) has also been used to change schedules in real time using predicted electricity prices [7]. Another significant factor to think about is uncertainty. To deal with prices that change and demand that isn't always predictable, strong optimization strategies have been suggested. These includes chance-constrained formulations [8], robust cost models utilizing L2-norms [9], and methodologies for worst-case scenario preparedness [10, 11]. These methods (discussed in Chapter 2) provide a strong mathematical foundation for robust solutions by focusing on convex optimization and stochastic programming.

Environmental Perspectives

Not every source of electricity are actually clean, and the volume of carbon they make can change depending on the time of day and the mix of energy sources. Charging electric vehicles (EVs) when the grid gets its power from renewable sources can cut emissions by a lot. Several studies combine goals for the economy and the environment. For example, stochastic unit commitment models [1] and two stage stochastic formulations for solar powered charging stations [2] have been developed. More advanced frameworks even use both stochastic and robust optimization to find a balance between cutting costs and cutting emissions [12].

Towards Practical Tools and Digital Twins

Most real charging stations still use simple rules like FIFS, even though these changes have been developed. This shows the difference between research and everyday life. Digital twins have been suggested as a way to connect both of these factors. They give operators and researchers a virtual copy of the charging system so they can try out different strategies before using them in the real world [13, 14].

Since they allow for complicated grid interactions, integrate real data, and record user behavior, digital twins are incredibly powerful tools. They are therefore ideal to evaluate optimization models in practical settings.

1.3 Scope of This Thesis

This thesis implements and compares three smart charging optimization models for electric vehicle (EV) infrastructure: the *Power Allocation Model* [9], the *Emission CO_2 Model* [15], and the *Cohort Model* [16]. Every model analyzes at a different facet of the EV charging problem, such as the cost, the environmental impact, and recommendations for improving smart charging methods.

These models are implemented using a digital twin framework created for this project via Python. This simulation platform integrates real electricity price data from Italy with synthetic EV arrival scenarios to assess each model's performance under controlled and repeatable conditions [17]. Researchers and decision-makers can modify crucial parameters like emission intensity, power capacity, number of EVs, and time horizon due to the framework's modular design.

Our work in this thesis includes two key goals. Evaluating each model's performance in terms of cost, emissions, and user satisfaction is the first step. The second is to show how digital twins can be used to test and simulate smart charging strategies before they are put into practice in real systems. By analyzing the models in both theoretical and simulated environments, the work demonstrates the differences in complexity, scalability, and performance. A flexible data-driven framework that can speed up the shift to smarter and greener EV charging infrastructure is presented in the thesis' conclusion.

Chapter 2

Preliminaries

In this chapter, we briefly introduce the mathematical basis needed for the following of this thesis. Concepts like convex optimization theory, Monte Carlo simulations, robust optimization under uncertainties, are very useful when dealing with these kind of smart charging problems.

2.1 Optimization Models

Generally speaking, optimization is useful because it allows us to convert multiple goals such as cutting emissions, energy expenses, or customer wait times into exact formulas. Following that, these formulations are resolved within the realistic parameters defined by the operational and technical limitations of the system.

A basic deterministic optimization problem can be written as:

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & f(x) \\ \text{s.t.} \quad & g_i(x) \leq 0 \quad \forall i \in \{1, \dots, m\}, \\ & h_j(x) = 0 \quad \forall j \in \{1, \dots, p\}, \end{aligned}$$

where $f(x)$ is the objective function to be minimized (usually related to some cost function), $g_i(x)$ are inequality constraints (like power or resource limits), and $h_j(x)$ are equality constraints (e.g., balance constraints or conservation laws).

2.2 Convex Optimization

The field of mathematics known as convex optimization studies problems in which the feasible region and the objective function share a structural characteristic called convexity. Convex problems are more predictable than general optimization problems, which can have unstable solutions or multiple local minima. For real-world applications in engineering, economics, machine learning, control systems, and operations research, they enable effective algorithms with provable convergence guarantees [18, 19].

Formally, a convex optimization problem takes the form:

$$\min_{x \in \mathcal{X}} \; f(x),$$

where $f : \mathbb{R}^n \to \mathbb{R}$ is a convex function, and $\mathcal{X} \subseteq \mathbb{R}^n$ is a convex set. A function f is convex if for all $x, y \in \mathcal{X}$ and $\gamma \in [0,1]$, it satisfies:

$$f(\gamma x + (1 - \gamma)y) \leq \gamma f(x) + (1 - \gamma)f(y).$$

This implies that the graph of f lies below the chord connecting any two points, ensuring that local minima are also global.

One of the good features that relates to convex optimization involves duality theory. For a primal problem, a lower bound on the optimal value provides a corresponding dual problem. Strong duality does hold under Slater's condition. This means that the primal and dual problems have optimal values that coincide. This property forms the basis for many algorithms as primal-dual interior point methods [18], plus aids in theoretical analysis.

Convex problems frequently appear in machine learning. Algorithms such as support vector machines (SVMs), logistic regression, ridge regression, and LASSO are formulated as convex optimization problems:

$$\min_{x \in \mathbb{R}^n} \sum_{i=1}^{m} f_i(x) + \lambda R(x),$$

where $f_i(x)$ measures prediction error, $R(x)$ is a convex regularization term (e.g., ℓ_1 or ℓ_2 norm), and λ balances the trade off [19].

For large scale or high-dimensional problems, first order methods like gradient descent and subgradient descent are widely used. These are especially effective in black box settings where only noisy gradient information is available. When curvature information is accessible, second order methods such as Newton's method or interior point methods can offer much faster convergence [20].

As dimensionality grows, non Euclidean techniques become valuable. Mirror descent, for example, generalizes gradient descent by incorporating a distance generating function ψ, leading to the iteration:

$$x_{k+1} = \arg\min_{x \in \mathcal{X}} \left\{ \langle \nabla f(x_k), x \rangle + \frac{1}{\eta_k} D_\psi(x, x_k) \right\},$$

where $D_\psi(\cdot, \cdot)$ is the Bregman divergence induced by ψ. Similarly, dual averaging methods aggregate gradients over time and adjust updates accordingly. These algorithms are especially useful in sparse learning, constrained optimization, and distributed systems [19].

Even when the original problem is non convex, convex optimization plays a critical role through relaxation. In these cases, constraints or objectives are approximated by convex counterparts. The relaxed problem is easier to solve, and the resulting solution can be rounded or refined post optimization. This approach is common in integer programming, control, and power systems optimization [18].

We can use convex optimization in many applications. For instance, Google's AdWords system allocates ad slots under budgetary constraints by solving a large-scale convex optimization problem [21].Effective and scalable power scheduling has been made possible in electrical grid operation by semidefinite relaxations of Optimal Power Flow (OPF [22]). Convex programs are used by model predictive control (MPC) frameworks in autonomous vehicles to plan trajectories in real time [23]. Convex resource allocation in wireless networks guarantees effective spectrum use while meeting QoS requirements [24]. Sparse recovery in MRI and compressed sensing has been revolutionized in medical imaging by total variation minimization [25].

In conclusion, convex optimization is basically a method of problem solving with a solid basis in math. It is beneficial as it is not only theoretical but also successfully utilized in real world situations. That makes it practical in fields where decisions depend on data, resources are limited, or safety is important.

Convex Relaxations and Tractability

Many practical problems are non convex due to integer constraints or nonlinearities. In such cases, *convex relaxations* are employed to approximate the original problem while preserving tractability. For instance, binary decision variables can be relaxed to continuous ones in [0,1], and non convex feasible regions can be over approximated by convex hulls [18].

A convex relaxation of a mixed-integer problem may take the form:

$$\min_{x \in \mathbb{R}^n} \quad f(x) \quad \text{s.t.} \quad x \in \text{conv}(\mathcal{X}_{\text{int}}),$$

where $\text{conv}(\cdot)$ denotes the convex hull of the original feasible integer set.

Such relaxations are common in large scale EV charging problems where real time decisions need to be computed efficiently. While the relaxed solution may not be strictly feasible, rounding and post processing can yield good practical solutions.

In summary, optimization models show an organized mathematical framework for decision-making in both uncertain and deterministic scenarios. Energy system stakeholders can create scalable and resilient control policies that strike a balance between performance, dependability, and cost effectiveness in complex environments by utilizing convex relaxation techniques, stochastic programming, and robust optimization.

2.3 Multi Objective Optimization

In real world EV charging scenarios, decisions are rarely based on a single goal. Instead, trade offs must be made between multiple, sometimes conflicting, objectives. For example, reducing electricity cost may lead to increased emissions if fossil fuel based energy is used, while minimizing emissions might require more expensive energy sources or additional delays.

Formulation

A typical multi objective optimization problem is written as:

$$\min_{x \in \mathcal{X}} \left[f_1(x), f_2(x), \ldots, f_k(x) \right],$$

where each function $f_i(x)$ represents a different objective (e.g., cost, emissions, waiting time). Since there is generally no single solution that minimizes all objectives simultaneously, the goal is to find Pareto optimal solutions.

Pareto Optimality

A solution x^* is said to be Pareto optimal if there is no other feasible solution that improves one objective without worsening another. The set of all such solutions forms the *Pareto front*, which helps decision makers understand the trade offs involved.

Weighted Sum and Scalarization

A common method to handle multiple objectives is to use a weighted sum:

$$\min_{x \in \mathcal{X}} \sum_{i=1}^{k} \lambda_i f_i(x),$$

where $\lambda_i \geq 0$ are weights reflecting the relative importance of each objective. In this thesis, such scalarization is used to combine cost and emissions using a trade-off parameter γ, as detailed in later chapters.

When creating EV charging plans that are both cost effective and ecologically friendly, multi objective optimization is essential. By adjusting weights or exploring the Pareto front, system operators can choose policies that best suit their specific goals [26, 27].

2.4 Sources of Uncertainty

The charging of electric vehicles (EVs) consists of a number of unknown and unavoidable variables that must be considered during simulation and optimization. These parameters will be explained below.

- **Arrival and Departure Times:**
  The behavior of EVs when they arrive and depart is one of the primary causes of uncertainty. Depending on their daily routines, charging habits, or work schedules, drivers may plug in their cars at different times. Both the energy demand and the scheduling flexibility are impacted by this fluctuation.
- **State of Charge and Energy Demand:**
  The amount of energy needed by each EV can vary significantly based on the initial state of charge (SoC) at arrival, battery capacity, and intended trip distance. This makes it difficult to predict how much energy each vehicle requires in advance.
- **Electricity Prices:**
  In systems with dynamic pricing, electricity costs change throughout the day based on supply, demand, and market conditions. These fluctuations introduce price uncertainty that directly affects the cost of charging.
- **Renewable Generation and Grid Conditions:**
  In grids that integrate solar or wind power, the available energy supply can be highly variable due to weather conditions. Similarly, grid overload or unexpected failures can impact the feasibility of charging decisions.

Incorporating these uncertainties into optimization models requires stochastic or robust techniques, as discussed in earlier sections. Accurate modeling of these factors leads to better decisions, especially in large-scale or real time systems [28, 29].

2.5 Stochastic Programming

Before talking about stochastic models, it's useful to compare them with deterministic programming. In deterministic optimization, all parameters such as arrival times, electricity prices, and energy demands are assumed to be known in advance. Stochastic programming, on the other hand, insert uncertainty directly into the model by treating unknown parameters as random variables with known probability distributions. This approach helps produce more reliable and flexible solutions under uncertain conditions [30, 31].

In many energy systems, key parameters such as electricity price, renewable generation, and user arrivals are uncertain. In these cases, *stochastic programming* is used to specifically include randomness into the model [30, 31]. The canonical two-stage stochastic programming problem is the following.

$$\min_{x \in X} \{f(x) + \mathbb{E}_\xi[Q(x, \xi)]\},$$

where: x is the first stage decision (e.g., scheduling EVs), ξ is a random vector capturing uncertainty (e.g., arrival time or price), $Q(x, \xi)$ is the optimal value of a second stage problem (the recourse function), defined as:

$$Q(x, \xi) = \min_{y \in Y(x,\xi)} q^\top y,$$

with y being second stage decisions (e.g., actual energy allocation), and $Y(x, \xi)$ defining feasible decisions under scenario ξ.

When a closed form of the expectation is unavailable, the *Sample Average Approximation (SAA)* is used [32]:

$$\min_{x \in X} \left\{ f(x) + \frac{1}{N} \sum_{i=1}^{N} Q(x, \xi^{(i)}) \right\},$$

where $\{\xi^{(1)}, \dots, \xi^{(N)}\}$ is a set of N sampled scenarios from the distribution of ξ. As $N \to \infty$, the solution converges almost surely to the true optimum.

2.6 Robust Optimization

While in stochastic programming we model uncertainties assuming to know their probability distributions, the *robust optimization* approach does not require any knowledge of them but instead optimizes against worst-case realizations within a predefined uncertainty set $\mathcal{U}$. The robust counterpart of a problem becomes [33]:

$$\min_{x \in X} \max_{\xi \in \mathcal{U}} f(x, \xi),$$

where $\mathcal{U}$ can be a box, ellipsoid, or polyhedral set. This formulation ensures performance under all possible values of $\xi \in \mathcal{U}$, providing a conservative but highly reliable solution.

For example, in energy systems where price volatility is significant (like the electricity price in smart charging), one may model uncertainty in prices π within bounds:

$$\pi \in [\bar{\pi} - \delta,\ \bar{\pi} + \delta],$$

and optimize:

$$\min_{x \in X} \max_{\pi \in [\bar{\pi} - \delta,\ \bar{\pi} + \delta]} \pi^\top x.$$

2.7 Monte Carlo Simulations

A popular numerical method for simulating and examining complex systems with uncertainty or unpredictable behavior is Monte Carlo simulation. Its fundamental idea is to estimate deterministic or stochastic quantities of interest through repeated random sampling. The method was first developed in the 1940s by scientists working on nuclear weapons research during the Manhattan Project and named after the gambling city of Monte Carlo to reflect its reliance on random outcomes [34].

Instead of relying on a single deterministic outcome, Monte Carlo simulation models thousands of possible system realizations to produce a more complete picture. This approach enables us to examine the probability of different outcomes, which aids in our understanding and management of the associated uncertainty. For example, just as repeatedly flipping a coin yields a better estimate of the probability of heads, increasing the number of simulations produces more accurate statistical predictions. Because of its intuitive appeal, Monte Carlo simulation is especially useful when system behavior is too complex or unpredictable to be modeled analytically [35].

Monte Carlo simulation can be used to generate multiple estimates of uncertain inputs, like electricity prices and EV arrival times, in the context of energy systems and smart charging. The effectiveness of various charging techniques in various life situations can then be evaluated using these insights.

Types of Monte Carlo Simulations

Monte Carlo techniques can be broadly classified into:

- **Static vs. Dynamic:** Static simulations focus on systems without time dependencies, while dynamic simulations (often called discrete-event simulations) model systems evolving over time.

- **Deterministic vs. Stochastic outputs:** Some simulations produce fixed outputs per input (e.g., Monte Carlo integration), while others include randomness throughout the process (e.g., agent-based energy demand simulations).

Monte Carlo Integration and Estimation

Mathematically, suppose we want to evaluate an expectation of the form:

$$\mu = \mathbb{E}[f(\xi)] = \int f(\xi)p(\xi)d\xi,$$

where ξ is a random variable with density $p(\xi)$, and $f(\xi)$ is the function of interest. Monte Carlo approximates this expectation using N independent samples $\{\xi^{(i)}\}_{i=1}^N$:

$$\hat{\mu}_N = \frac{1}{N}\sum_{i=1}^{N} f(\xi^{(i)}).$$

By the law of large numbers, this estimate converges almost surely to μ as $N \to \infty$, and the central limit theorem ensures that the error decreases at a rate of $\mathcal{O}(1/\sqrt{N})$.

Sample Average Approximation (SAA)

In stochastic optimization, expectations often appear in the objective function or constraints. The Sample Average Approximation (SAA), which provides a sample based average for the true expectation, is a popular method for resolving such issues. For a problem of the form:

$$\min_{x \in X} \left\{ f(x) + \mathbb{E}_\xi[Q(x, \xi)] \right\},$$

the SAA formulation becomes:

$$\min_{x \in X} \left\{ f(x) + \frac{1}{N}\sum_{i=1}^{N} Q(x, \xi^{(i)}) \right\},$$

where $\xi^{(i)} \sim p(\xi)$ are i.i.d. samples. As N increases, the SAA solution converges almost surely to the solution of the true problem under mild conditions [36].

Advanced Techniques: Variance Reduction and MLMC

While simple Monte Carlo estimation is unbiased, its convergence rate may be slow. To improve efficiency, variance reduction techniques are used:

- **Control Variates:** Leverages correlation with known variables to reduce variance.
- **Antithetic Sampling:** Uses pairs of negatively correlated samples.
- **Importance Sampling:** Samples more frequently from important regions of the input space.

Another advanced method is *Multilevel Monte Carlo (MLMC)*, which reduces computational cost by combining simulations at different levels of resolution. The idea is to simulate many inexpensive low-fidelity samples and only a few high-fidelity ones, balancing cost and accuracy [37].

Applications in the Real World

Monte Carlo methods have been applied in a range of domains:

- **Finance:**
  They are used to evaluate derivative pricing under uncertain market dynamics, model credit risk, and optimize portfolios by simulating thousands of future asset price paths [38].

- **Energy Systems:**
  Operators simulate probabilistic grid conditions to evaluate reliability, optimize power flow under renewable variability, and assess blackout risk [37].

- **Reliability Engineering:**
  Monte Carlo is used to estimate the failure probability of mechanical parts. For example, in the automotive industry, it models fatigue in vehicle components subjected to variable load cycles [39].

- **Logistics:**
  Supply chain models use stochastic sampling to simulate demand fluctuations, supplier delays, and stockout probabilities, helping firms optimize inventory policies under uncertainty [40].

Strengths and Limitations

Monte Carlo methods are appreciated for their:

- Simplicity and ease of implementation,
- Flexibility in handling non-linear, non-convex, or high-dimensional problems,
- Compatibility with black-box models and simulation-based systems.

However, they also come with drawbacks:

- Require large sample sizes to achieve low error,
- May be computationally expensive for real-time applications,
- Accuracy depends on quality and independence of samples.

In conclusion, Monte Carlo simulation provides a useful and efficient way to deal with uncertainty. It is a fundamental component of simulation-based approaches, which are widely utilized in both academic research and industrial practice due to its flexibility and ability to handle complex systems.

2.8 What is a heuristic?

Not all charging strategies are built from optimization models. In practice, many depots and public stations rely on simple rules *heuristics* that are easy to implement and understand. In this thesis, we use *First In First Served (FIFS)* as our main heuristic baseline. A heuristic is a rule based policy that makes decisions quickly without solving a mathematical program. Heuristics are useful as:

- **Baselines** to compare against optimized solutions,
- **Fast fallbacks** when computation or data is limited,
- **Explainable policies** that operators can implement immediately.

FIFS policy (queue discipline)

FIFS assigns charging power to vehicles strictly by arrival order. When power or sockets are limited, early arrivals are served first; later arrivals wait.

Notation. Let $\mathcal{Q}(t)$ be the queue of connected electric vehicles ordered by arrival time, P_t the total available power at time t, $\bar{p}$ the limit of the per EV rate, and S_t the number of sockets (if applicable).

Rule. At each time step t:

1. Serve the first $\min\{|\mathcal{Q}(t)|, S_t\}$ vehicles in the queue.
2. Allocate to each served EV up to $\min\{\bar{p}, P_t/\text{served}\}$.
3. Update the queue by removing EVs that become fully charged or depart.

Why include FIFS?

FIFS reflects how many real-world sites operate today. It is:

- **Simple**: no forecasts, no solvers.
- **Predictable**: operators and users understand the rule.
- **Computationally light**: $\mathcal{O}(N)$ per step with a queue.

However, FIFS ignores price and emissions signals and may be *unfair* to later arrivals, which can increase cost and occasionally leave users undercharged at departure.

Cohort-style FIFS

For comparison with our Cohort Model, we also use *cohort style FIFS*: EVs are grouped into energy demand classes and the service proceeds class by class in arrival order. This aggregated version keeps the spirit of FIFS while matching the state representation used in Chapter 3.

When to use a heuristic

Heuristics are appropriate when:

- data is scarce or forecasts are unreliable,
- strict response-time constraints preclude optimization,
- an interpretable, "no surprises" policy is preferred.

In this thesis, FIFS serves as a transparent baseline to highlight the gains achieved by optimization in terms of cost, emissions, and user satisfaction.

Chapter 3

Models

Building on the optimization concepts introduced in Chapter 2, this chapter presents the three main mathematical models developed and implemented in this thesis, together with a simple baseline strategy (FIFS). Each model addresses smart charging of electric vehicles (EVs) from a different angle. The models considered are the following:

- *Power Allocation Model* [41]. The main goal of this model is to minimize the overall electricity cost, while at the same time taking into account the dissatisfaction of users if their charging demand is not fully met.
- *Emission CO_2 Model* [15]. Here, the emphasis shifts to environmental aspects: the optimization explicitly includes CO_2 emissions in the objective, providing a trade off between cost and sustainability.
- *Cohort Model* [16].The model groups EVs into classes that reflect their charging requirements and availability, and tracks how these groups shift from one state to another as time progresses.
- *First In First Served (FIFS) Model*. This is a straightforward reference strategy in which vehicles are charged in the order of their arrival until they are either fully charged or leave the system. Although FIFS is not an optimization problem, it is useful as a baseline for comparison. To make the comparison fair with the Cohort Model, we implement a cohort style version of FIFS: EVs are divided into discrete classes, their transitions are tracked over time, and charging power is allocated sequentially following arrival order. While clearly less efficient than optimization based strategies, this heuristic reflects how many real world charging stations actually operate.

Together, these models show different ways to manage EV charging: lowering costs, reducing environmental impact, and improving fairness and efficiency across the fleet. Despite their different perspectives, all models share the same mathematical backbone of convex optimization, which allows them to be solved reliably with standard numerical solvers and compared consistently.

Each formulation translates the smart charging problem into a set of mathematical components:

- **Objective functions** that capture the goals of the system whether minimizing cost, emissions, or energy deficit;
- **Decision variables** representing charging allocations or fleet states;
- **Constraints** that reflect physical limitations, such as charger capacity, vehicle availability, or power budgets.

Despite these shared foundations, the models differ significantly in how they handle system complexity and uncertainty. For example, the Power Allocation and Emission CO_2 models operate at the individual vehicle level with explicit constraints per EV, whereas the Cohort Model adopts a more aggregated, class based approach to capture the collective dynamics of a fleet. Moreover, models vary in how they address uncertainty: deterministic parameters can be replaced with scenario based inputs using Monte Carlo simulations (as described later), enabling realistic stress tests under price volatility, fluctuating arrivals, and changing emission intensities.

Each model changes theory into helpful instruments to manage EV charging in actual operating environments by building on the optimization concepts previously discussed, such as convexity, dual objectives, and stochastic methods. Each model's formulation, underlying presumptions, and decision-making guidelines are thoroughly explained in the remaining portion of the chapter.

3.1 Power allocation Model

This model is formulated as a classical convex optimization problem that integrates both economic and service quality considerations. Its primary goal is to minimize the total electricity cost incurred by charging a fleet of electric vehicles, using real or simulated electricity price profiles. However, unlike traditional cost minimization approaches, this model also explicitly accounts for customer dissatisfaction by penalizing any shortfall in the energy delivered to each EV.

In practice, energy shortfalls may occur when available power is limited or when vehicles are connected for a shorter period than required to meet full charging needs. To reflect this, the model introduces a decision variable that captures the energy *deficit* per vehicle, the difference between the requested energy and the energy actually received during the stay of the vehicle.

The objective function combines two components:

- The first term captures the monetary cost of energy drawn from the grid, calculated by multiplying the allocated power by time varying electricity prices;
- The second term introduces a dissatisfaction penalty, linearly proportional to the energy deficit, and scaled by a parameter γ that reflects how much importance the operator assigns to meeting user demand.

By tuning the weight γ, one can balance the trade off between minimizing cost and ensuring high service quality. A low γ places more emphasis on reducing grid expenses, possibly at the cost of user satisfaction, while a higher γ prioritizes full energy delivery even if it means accepting higher electricity prices.

Because all variables and constraints in this formulation are linear, and the objective is convex, the model can be solved efficiently using standard linear programming solvers. Its simplicity makes the model practical for real time use or large scale simulations where fast solutions are needed. Additionally, it provides a point of comparison with advanced algorithms which consider dynamic or environmental factors, that are covered in the following sections.

Objective and Structure

The decision variable Y denotes the charging power assigned over time to each electric vehicle, and e captures the energy deficit (the amount of energy requested but not delivered). The objective includes:

- the total cost of charging, computed via the price vector π, and
- a dissatisfaction penalty scaled by a tunable weight γ, which reflects the trade off between minimizing cost and meeting energy demand.

$$\begin{aligned}
\min_{Y,e} \quad & \pi^\top \delta Y \mathbf{1} + \gamma \mathbf{1}^\top e \\
\text{s.t.} \quad & \delta Y \mathbf{1} \geq L - e, \\
& Y \mathbf{1} \leq C, \\
& Y \geq 0, \quad Y \leq \bar{p} \mathbf{1}, \quad e \geq 0, \\
& Y_{ti} = 0, \quad \forall t, i : A_{ti} = 0.
\end{aligned} \tag{3.1}$$

The parameters are:

- $L \in \mathbb{R}^N$: energy demand, with L_i representing how much energy EV i wants to receive during its stay.
- $C \in \mathbb{R}^T$: power capacity of the grid or charging station at each time step.

- $\bar{p}$: Europe's most widely used charging ports limit
- $A \in \{0,1\}^{T \times N}$: availability matrix, where $A_{ti} = 1$ indicates that vehicle i is connected to the charger at time t.

Using known or predicted values for availability and price, this formulation takes a deterministic approach. The problem can be solved quickly and reliably because it is convex. A technique for examining various price scenarios and evaluating the precision of the findings is provided by Monte Carlo sampling.

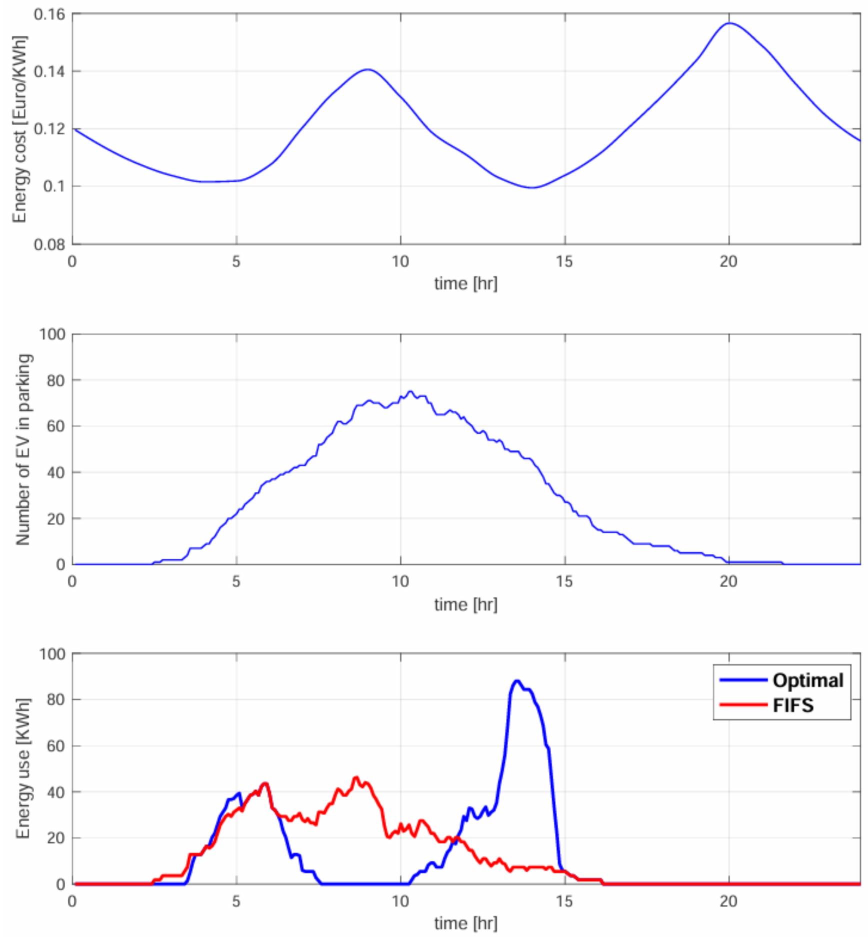

Figure 3.1: Illustration of the Power Allocation model under historical price data [9]. Top: electricity price profile; Middle: EV parking occupancy; Bottom: comparison of energy allocation between the optimized strategy and the FIFS baseline.

Figure 3.1 illustrates how the Power Allocation model operates. The top plot shows daily electricity prices, the middle plot shows the number of EVs parked, and the bottom plot compares the optimized schedule with the naive FIFS approach. The optimized model shifts charging to cheaper hours and avoids peaks, whereas FIFS ignores price signals, resulting in higher costs.

3.2 Emission CO_2 Model

This model looks not only at the cost of charging but also at its environmental impact. Alongside lowering electricity expenses, it considers the carbon emissions from grid power.These elements support a more economical and ecologically responsible method of energy management. If we look for decarbonization goals and develop more flexible and cleaner power systems, this kind of strategy is essential. In practice, the model is set up as deterministic, with electricity prices and carbon emission rates treated as known beforehand. However, these inputs are allowed to vary dynamically over time, enabling the model to adapt to changing grid conditions such as increased renewable energy penetration or peak emission periods. In this sense, the model integrates time sensitive environmental signals, acknowledging that the carbon intensity of electricity is not constant throughout the day.

Mathematically, the objective function is the weighted sum of two terms:

- The total electricity cost, computed by multiplying the amount of energy allocated at each time step by the corresponding electricity price; and
- The total CO_2 emissions, calculated by using time varying emission factors that reflect the environmental impact of electricity consumption.

The weighting parameter λ allows for a flexible trade off control between economic and environmental goals. A value of $\lambda = 0$ reduces the model to pure cost minimization, while higher values shift the focus to emission reduction. System operators, energy planners, or charging station managers can match optimization to their sustainability goals or policy objectives thanks to this flexibility.

Figure 3.2 shows how the Emission CO_2 Model operates. The top plot compares the price of electricity and the cost of CO_2 emission over time, highlighting that cheaper energy is not always cleaner. The bottom plot illustrates how the optimized model shifts charging to periods with lower prices and emissions, unlike the FIFS strategy, which ignores these signals. This demonstrates the value of combining cost and environmental factors in the decision making process.

Model Motivation

In this formulation the model balances electricity cost with the carbon footprint associated with charging. It is particularly useful when the electricity supply is mixed.

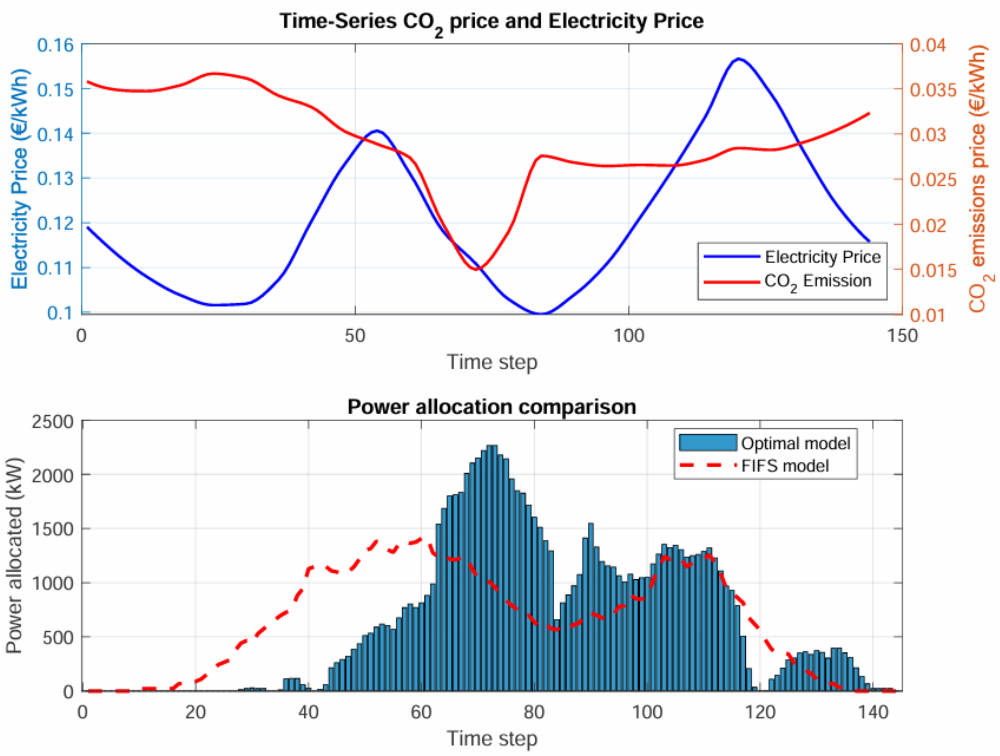

Figure 3.2: Example of the Emission CO_2 Model in action. electricity price and CO_2 emission cost compare the optimized model (blue) with the FIFS strategy (red). [15]

Mathematical Formulation

$$\min_{Y}\left[\sum_{t=1}^{T}\pi_t\left(\sum_{i:t\in[a_i,d_i]}\delta Y_{ti}\right)+\lambda\sum_{t=1}^{T}\pi_t^{CO_2}\left(\sum_{i:t\in[a_i,d_i]}\delta Y_{ti}\right)\right] \tag{3.2}$$

$$\begin{aligned}\text{s.t.}\quad & \sum_{t=a_i}^{d_i}\delta Y_{ti}\geq L_i,\quad \forall i,\\ & \sum_{i=1}^{N}Y_{ti}\leq C_t,\quad \forall t,\\ & 0\leq Y_{ti}\leq\bar{p},\quad \forall i,t,\\ & Y_{ti}=0,\quad \forall t\notin[a_i,d_i],\forall i.\end{aligned}$$

The parameters are as follows.

- $L\in\mathbb{R}^N$: energy demand, with L_i representing how much energy EV i wants to receive during its stay.

- $C \in \mathbb{R}^T$: power capacity of the grid or charging station at each time step.
- $\bar{p}$: maximum charging rate per EV (kW), determined by the station limits.
- $A \in \{0,1\}^{T \times N}$: availability matrix, where $A_{ti} = 1$ indicates that vehicle i is connected to the charger at time t.
- $Y \in \mathbb{R}^{T \times N}$: the charging power matrix, where Y_{ti} denotes the power allocated to vehicle i at time step t.
- $\pi_t^{CO_2}$: Emission adjusted electricity price at time t, derived from the emissions vector e_{CO_2} (in g/kWh) scaled by a carbon cost constant. Specifically,

$$\pi_t^{CO_2} = C_0 \cdot e_{CO_2}(t)$$

  where $C_0 = 70$ \$/ton $= 64$ $£/ton$, and in the Vietnamese context used in our study.

Discussion

The weighting parameter λ is the most crucial component of the Emission CO_2 Model. The operator can choose whether to prioritize more on cost reduction or emissions by adjusting this parameter, which functions as a slider. For instance, if

Since λ is small, the optimizer's primary goal is to maintain a low electricity bill. The model prioritizes charging when the grid is cleaner if λ is large, even if it means paying a marginally higher price. In a way, the model provides flexibility to different operational plans or goals.

The model also makes the assumption that the production and consumption of electricity are separate. The carbon intensity of electricity, represented by $\pi_t^{CO_2}$, is taken as an external input that varies over time. This reflects the fact that the emissions linked to grid electricity are not constant during the day they change with the mix of renewable and fossil fuel generation.

From a mathematical perspective, the formulation remains linear and convex, which means it can be solved quickly and reliably even for large systems. This makes it practical for real time applications or large scale fleet simulations.

Key advantages:

- Promotes charging at times when the grid is cleaner, leading to lower emissions.
- Can be adapted to different seasons, regions, or future changes in emission profiles.
- Provides a balance between economic savings and environmental sustainability, depending on the chosen value of λ.

3.3 Cohort Model

The *Cohort Model* is inspired by state based control frameworks and offers a scalable, system level representation of electric vehicle (EV) charging dynamics. Rather than managing each EV individually which becomes computationally intensive as the fleet grows this model adopts a class based approach. Vehicles are grouped into *cohorts* or *charging classes* based on their remaining energy demand. From almost empty batteries to fully charged ones, each class stands for a step in the charging process.

This abstraction shifts the focus from micro level decisions (per EV charging profiles) to macro level behavior (how many EVs are in each state and how they transition between states over time). It is especially helpful for centralized management in huge public or commercial charging infrastructures because it shows the overall dynamics of an entire EV fleet.

The mathematical formulation that supports it is dynamic. The system evolves according to a discrete time state transition model, where the number of vehicles in each class at time $t + 1$ depends on the current state $x(t)$, arrivals $a(t)$, and control decisions $c(t)$ (i.e., how many vehicles to charge in each class). Transition matrices $A(t)$ and B govern how vehicles naturally progress due to departures and how charging actions impact state shifts.

Goals for service quality and the economy are combined in the objective function. It includes:

- The total electricity cost incurred over time, weighted by dynamic price signals
- A dissatisfaction penalty based on the number of vehicles remaining in lower charge states, weighted by class specific coefficients.

Important aspects of fleet behavior in the real world are captured by this model:

- It enables prioritization of undercharged vehicles through the dissatisfaction weights.
- It accommodates dynamic arrivals and departures through $a(t)$ and $A(t)$.
- It naturally scales to fleets of arbitrary size by adjusting the number of classes.

The Cohort model can find moderation between efficiency and accuracy. Because it reduces the optimization problem while preserving key charging dynamics, it is ideal for online control or long-term planning. Because Monte Carlo methods can be used to sample uncertain arrivals, the model also easily integrates with stochastic programming techniques (Section 2.1.1). So, the Cohort Model is a practical replacement for individual EV scheduling. Its state transition foundation

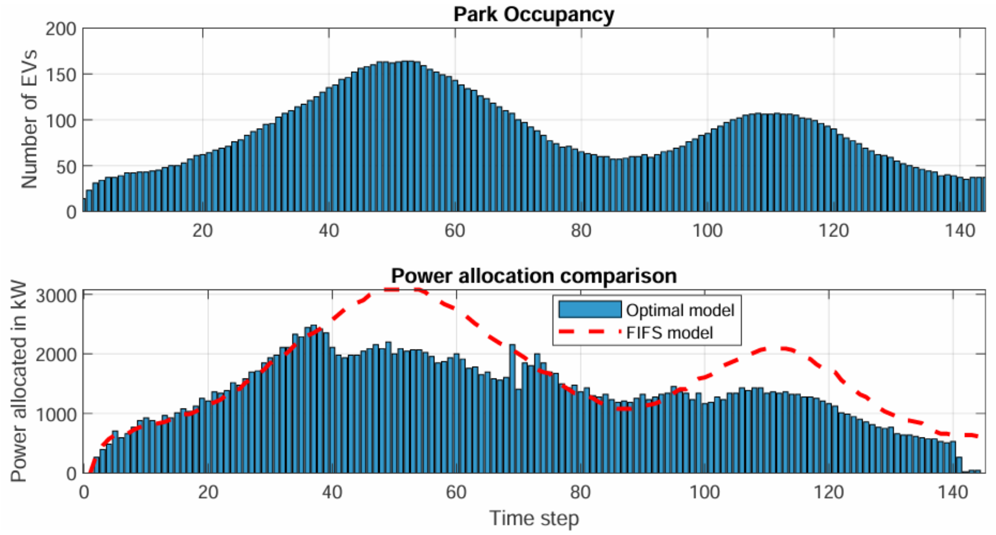

Figure 3.3: Cohort Model- park occupancy over time, power allocation comparison between the optimized cohort schedule and the FIFS baseline.. [16]

makes it ideal for capturing changing system behavior under uncertainty, and its class-based structure lowers dimensionality while maintaining accuracy.

Figure 3.3 shows how the cohort model uses class based states to track fleet evolution (top) and schedule charging (bottom). As occupancy rises, the optimized policy spreads charging over time and avoids high peaks, while the FIFS baseline closely follows arrivals and creates larger bursts of power. This illustrates why grouping EVs into classes enables scalable, smoother control at the fleet level.

System Representation

The model uses the state vector $x(t)$ to track the number of EVs in each class at time t, and $c(t)$ as the number of vehicles receiving charging in each class. Transition dynamics are captured using a matrix model, and arrivals are included via $a(t)$. The model minimizes both cost and dissatisfaction using the this formulation:

$$\begin{aligned} \min_{C,X} \quad & \sum_{t=0}^{T-1} \pi_t \Delta P_0 \mathbf{1}^\top c(t) + \gamma \sum_{t=0}^{T-1} \beta^\top x(t) \\ \text{s.t.} \quad & 0 \le P_0 \mathbf{1}^\top c(t) \le \bar{P}, \quad \forall t, \\ & 0 \le c(t) \le x(t), \quad \forall t, \\ & x(t+1) = A(t)x(t) + a(t) + Bc(t), \quad \forall t. \end{aligned} \tag{3.3}$$

Parameters

The parameters for the Cohort Model are as follows:

- $x(t) \in \mathbb{R}^{n+1}$: state vector at time t, where each entry represents the number of EVs in a given energy class (grouped by remaining energy to be charged).
- $c(t) \in \mathbb{R}^{n+1}$: control vector/matrix at time t, indicating the number of EVs from each class being charged during that time step.
- $a(t) \in \mathbb{R}^{n+1}$: arrival vector/matrix, denoting the number of new EVs entering each class at time t. It is generated via probabilistic models (Poisson).
- $A(t) \in \mathbb{R}^{(n+1)\times(n+1)}$: departure transition matrix at time t, modeling the proportion of EVs that remain or depart from each class.
- $B \in \mathbb{R}^{(n+1)\times(n+1)}$: charging transition matrix, describing how the control action $c(t)$ shifts vehicles from one class to another as energy is supplied.
- $\beta \in \mathbb{R}^{n+1}$: dissatisfaction weight vector, assigning higher penalty to vehicles with higher remaining demand (i.e., lower charging class).
- $\pi_t \in \mathbb{R}$: electricity price at time t (€/kWh), used to compute the cost of charging over time.
- $\Delta \in \mathbb{R}$: duration of each time step (minutes), used to convert power to energy.
- $P_0 \in \mathbb{R}$: fixed charging power allocated per vehicle (kW), assuming uniform charging rates across all classes.
- $\bar{P} \in \mathbb{R}$: total power limit of the charging station (kW), representing the maximum grid capacity at any time.

The framework of stochastic programming naturally supports this formulation. Because it is convex and allows for flexible arrival modeling using Monte Carlo techniques, the method is scalable even as the state space grows.

Summary

All three models are grounded in the optimization frameworks introduced in Chapter 2, but each takes a unique approach to smart EV charging:

- **Power Allocation Model:** This model captures the trade off between electricity cost and user dissatisfaction due to unmet energy demand. *Advantages:* It is simple and fast to solve, balances cost and service quality effectively, and is well suited for individual EV tracking. *Limitations:* It does not account for emissions, scales poorly for very large fleets, and assumes fixed vehicle availability.

- **Emission CO_2 Model:** This model extends the objective function to incorporate environmental costs. *Advantages:* It explicitly includes environmental objectives, leverages real price and emissions data, and provides flexibility through a tunable trade off parameter. *Limitations:* It assumes perfect knowledge of future conditions (deterministic), does not capture fleet dynamics, and its results are sensitive to the choice of emission weight.

- **Cohort Model:** This model adopts a state transition structure that groups EVs by energy class, allowing scalable fleet level scheduling under uncertainty. *Advantages:* It captures fleet evolution, is efficient for large populations, and naturally accommodates probabilistic arrivals and departures. *Limitations:* It requires discretization of energy states, is less detailed at the level of individual EVs, and needs more calibration for transition matrices (A, B).

These three models all present three distinct methods for EV smart charging: reducing individual costs, reducing environmental impact, and enhancing large scale coordination.Depending on the specific goals or preferences, each model can be used independently. In the following chapter, simulations that represent real world scenarios are used to evaluate their performance.

To provide a meaningful baseline for comparison, we also implemented a **First In First Served (FIFS)** strategy. Unlike the optimization-based models, FIFS simply allocates charging power in arrival order, ignoring prices, emissions, or demand forecasts. *Advantages:* It is very simple, predictable, and computationally light, making it easy to implement in real world systems. *Limitations:* It ignores market signals and emissions, can lead to unfairness for later arrivals, and often results in higher costs or undercharged users.

In the context of the Cohort Model, we extended FIFS into a class based version, applying the same arrival order logic within each group of vehicles sharing similar energy demand characteristics.

Chapter 4

Simulations and Results

This chapter presents the simulation setup to evaluates the three optimization models introduced in Chapter 3, selected parameters, and numerical results for the three models described in Chapter 3: the *Power Allocation Model*, the *Emission CO_2 Model*, and the *Cohort Model*. Here the focus is on testing how each model responds to different conditions such as shifts in electricity prices, differences in energy demand, and variations in EV arrivals. Each model is simulated over a 24 hour period. The results are evaluated based on key performance indicators like total cost, CO_2 emissions, user satisfaction (measured through unmet demand or dissatisfaction metrics).

While the general structure of the simulation remains consistent, the specific parameter values such as arrival distributions, power limits, or weights may differ between models to better reflect their individual assumptions and design goals. To manage this, a modular simulation framework was implemented in Python, allowing flexible configuration and easy switching. The chapter is organized in which each section looks at a single model, explaining its numerical analysis, simulation results, and parameter choices. The key results and insights are compiled in a final comparison and discussion.

4.1 Power Allocation Model Results

This section presents the simulation environment and results for the *Power Allocation Model*, based on the implementation of the optimization problem discussed in Chapter 3. Evaluating the model's capacity to allocate power efficiently, control charging expenses, and minimize customer dissatisfaction under actual charging situations is the purpose.

Simulation Setup

We simulate a typical day of electric vehicle charging over a 24 hour horizon, which is discretized into 288 time steps of 5 minutes each ($\delta = 5$ minutes). The scenario involves a fleet of $N = 100$ electric vehicles, representing a charging station of an appropriate size, like a commercial lot or workplace.

To reflect realistic user behavior, each EV's arrival time a_i is generated using a Beta distribution Beta($\alpha = 5.52, \beta = 12$), scaled to the 24 hour window. This produces a natural concentration of arrivals around 7:00 AM similar to what one might observe in morning commute scenarios. Departure times d_i are drawn from a normal distribution centered at 6 hours after arrival, with a standard deviation of 1 hour, ensuring that most vehicles remain parked for a typical workday.

The energy demand L_i of each vehicle is defined as a fraction of its theoretical maximum, based on its parking duration. Specifically, $L_i = w_i \cdot \bar{L}_i$, where w_i is randomly drawn from a uniform distribution between 0 and 1, and $\bar{L}_i = \bar{p}(d_i - a_i)$ represents the maximum energy that could be delivered to EV i during its stay. The maximum charging rate per vehicle is set to $\bar{p} = 22$ kW, and the charging station has a constant total power capacity of $C_t = 1130$ kW available at each time step.

The electricity price profile π_t used in the simulations is based on actual data from the Italian energy market. We used hourly averages collected over 10 years (2015-2024), then interpolated and resampled them to match the 5 minute resolution of our simulation.

To incorporate user satisfaction, we assign a dissatisfaction penalty $\gamma = 0.157$ €/kWh. This parameter encourages a fair balance between price and user satisfaction while guaranteeing that the optimizer is not entirely focused on reducing costs at the risk of overcharging vehicles.

Results and Discussion

Our simulations show the Power Allocation Model functions well in practical situations. The optimizer automatically shifts charging to less expensive hours, saving users money without asking them to make any adjustments. By finding a balance between energy consumption and the actual availability of vehicles, it also makes the allocation more realistic than with a naive approach. Another important aspect of the model is how close it was to grid limits, which guarantees that the system won't be overloaded during times of high demand. Together, these actions show that the Power Allocation Model can successfully find a balance between affordability, usefulness, and grid stability.

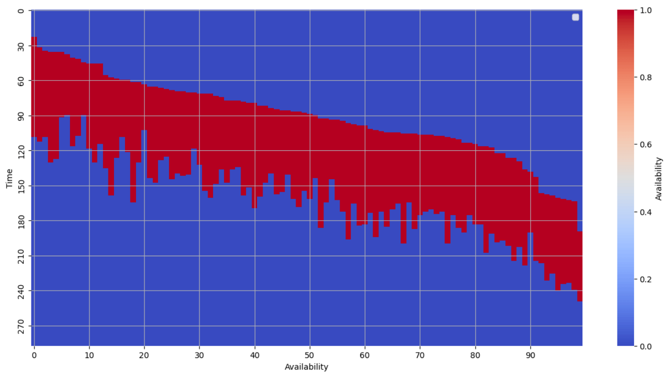

Figure 4.1: EV Availability Matrix – sorted by arrival time.

The availability matrix of the 100 EVs is shown in Figure 4.1, with the presence of vehicles in each time slot shown in darker (red) areas. There is a certain pattern: the majority of cars arrive early in the morning, filling the parking spots, and then slowly leave in the late afternoon. This is similar to how employees usually charge their cars at work, plugging them in at the beginning of the day and unplugging them at the end of their shift. Because it creates unexpected availability peaks that the model must control, this type of pattern is significant. There is a strong competitor for charging resources in the middle of the day when the lot is almost full. On the other hand, as cars depart in the evening and at night, demand falls rapidly. Since these dynamics have a direct impact on charging scheduling and the optimizer's ability to maintain a balance between cost reduction and user satisfaction, it is essential to capture them.

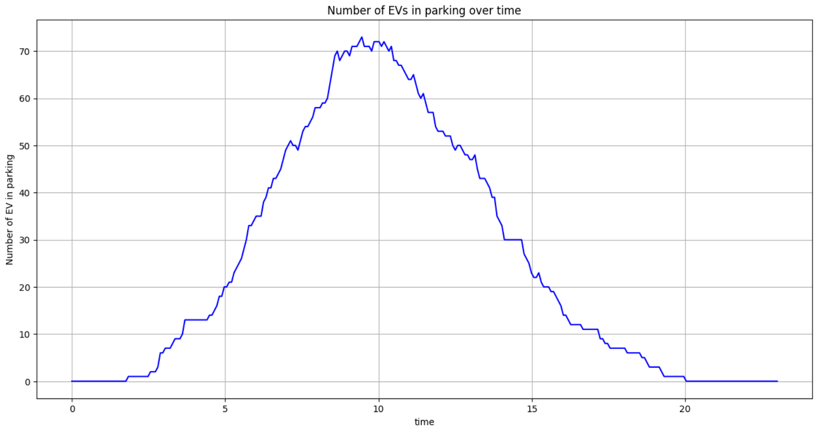

Figure 4.2: Number of EVs parked over time.

The number of electric cars in the parking lot changes during the day, shown in Figure 4.2. As cars arrive in the morning, the curve rises rapidly. It levels off around midday, when most of the fleet is parked, and then it gradually decreases as cars depart in the evening. Because cars are typically plugged in during working hours and unplugged before drivers leave for home, this produces a typical daily charging cycle that simulates workplace or campus settings.

The management of charging is directly affected by this pattern. The demand for charging rapidly increases as the lot fills up in the morning, putting pressure on the system to allocate power effectively. Since most cars are already connected during the midday peak, the optimizer can move loads to less expensive hours or times when the grid isn't under as much demand. There are fewer scheduling opportunities in the evening when cars disconnect, making any missing demand obvious. Since it affects the timing and the total amount of power that can be distributed, it is crucial to know this daily cycle.

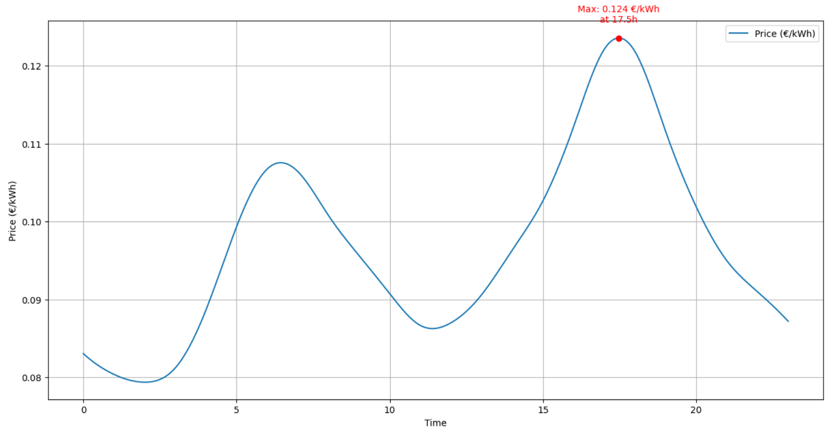

Figure 4.3: Electricity Price Profile (Resampled from Historical Data).

The electricity price profile used in the simulation is shown in Figure 4.3. It features two main peaks: a smaller one in the morning and a more pronounced spike in the late afternoon, consistent with typical residential and commercial load patterns in the Italian energy market.

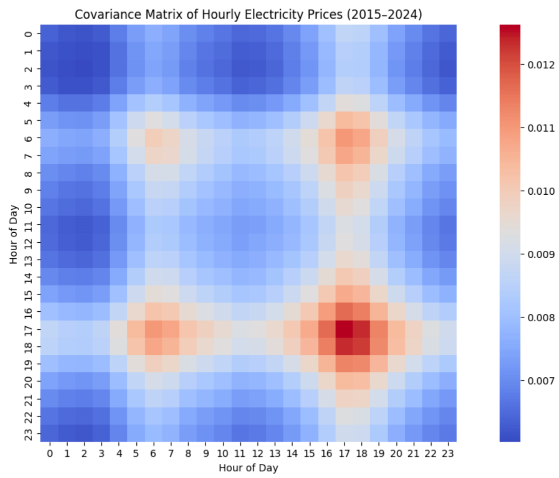

Figure 4.4: Hourly Covariance of Electricity Prices (2015–2024).

To capture how prices vary over time, Figure 4.4 shows the covariance matrix of hourly electricity prices based on 10 years of data. A covariance matrix is a compact way of describing how different variables move together. While evaluating different possible outcomes through scenario-based modeling or stochastic sampling, this information is helpful.In our case, however, we relied on the nominal price profile for the simulations.

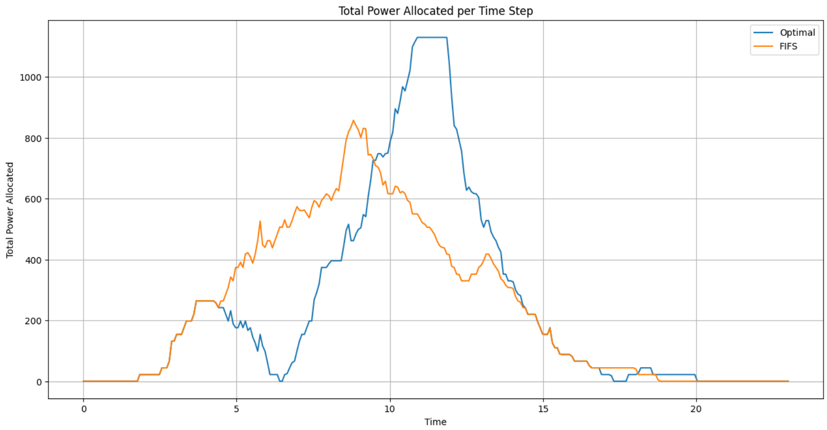

Figure 4.5: Total Power Allocated over Time – Optimal vs. FIFS.

It is clear in Figure 4.5 how the optimal model differs from the basic First-In-First-Served (FIFS) method. The optimal plan stops charging initially and then shifts the majority of charging to the late morning and early afternoon, when electricity is less expensive. By doing so, it profits on the price plateau, maintains cost reductions, and keeps the grid from experiencing too much stress at the wrong times.

In contrast, FIFS does not take into account price changes or grid conditions; instead, it charges cars as soon as they arrive. It means that energy is used less effectively, frequently during more costly hours, which raises the final cost. This side-by-side comparison shows why optimization is important: the optimal model listens to external signals and makes smarter choices that balance cost and system stability rather than adopting a simple "plug-and-charge" approach.

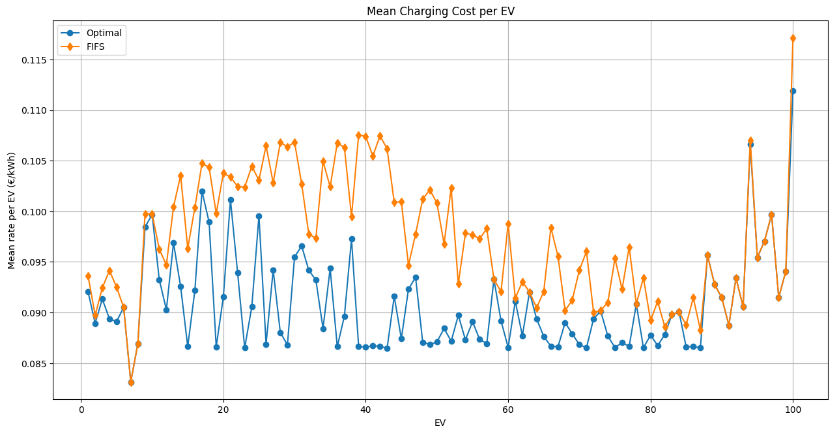

Figure 4.6: Mean Charging Cost per EV – Optimal vs. FIFS.

The average charging cost per EV is plotted in Figure 4.6, which makes this cost advantage more obvious. For almost all EVs, the best approach always results in lower or comparable costs, particularly during peak hours, while FIFS experiences more fluctuation and fluctuations.

The findings show that the Power Allocation Model efficiently meets user requirements while cutting expenses. It is significantly more efficient than the heuristic algorithm due to its sensitivity to price variation and its fairness adjustment (via γ). The total cost of electricity over the simulation horizon is one of the primary performance indicators. The Power Allocation Model generated a cost of 550.751 € by strategically allocating charging across less expensive time slots. The *First In First Served (FIFS)* policy, on the other hand, led to a higher cost of 575,048 € because it charges on arrival regardless of the price incurred.

The comparison illustrates the benefit of using price information in the optimization. Despite being used as the standard for uncontrolled systems, FIFS's leads to ineffective energy distribution, especially during peak hours.

4.2 Emission CO_2 Model Results

Simulation Setup

With a 24-hour horizon split into 144 time steps of 10 minutes each ($\delta = 10$ minutes), the simulation simulates an average day of EV charging. A total of $N = 300$ electric cars are modeled. $\bar{p} = 22$ kW is the fixed charging rate per EV. $C_t = N \cdot \bar{p}$ gives the station's total power capacity, which means it can support all connected EVs at once. To balance cost and emissions in the objective function, a trade-off coefficient $\lambda = 9$ is added.

The EVs are divided into two groups (morning and afternoon), each of which is modeled using a Beta distribution in order to produce realistic arrival and departure patterns. Arrivals peak in the morning at :00$AMandintheafternoonat$: 00 PM. Accordingly, the beta distribution is scaled:

$$a_i \sim \lfloor T \cdot \text{Beta}(\alpha_{\text{mor/after}}, 12) \rfloor$$

with:

$$\alpha_{\text{mor/after}} = \frac{1 + h_{\text{mor/after}} \cdot (\beta - 2)}{1 - h_{\text{mor/after}}}$$

where $h_{\text{mor}} = 7/24$ and $h_{\text{aft}} = 17/24$.

Departure times d_i are computed as:

$$d_i = \min(a_i + \tau, T)$$

where τ is drawn from a normal distribution, depending on morning or afternoon behavior: $\tau_{\text{mor}} \sim \mathcal{N}(7, 4^2)$ and $\tau_{\text{aft}} \sim \mathcal{N}(3, 1^2)$ (in hours).

Each EV's energy demand is computed as:

$$L_i = \beta(2, 2) \cdot \bar{p} \cdot (d_i - a_i)$$

Results and Discussion

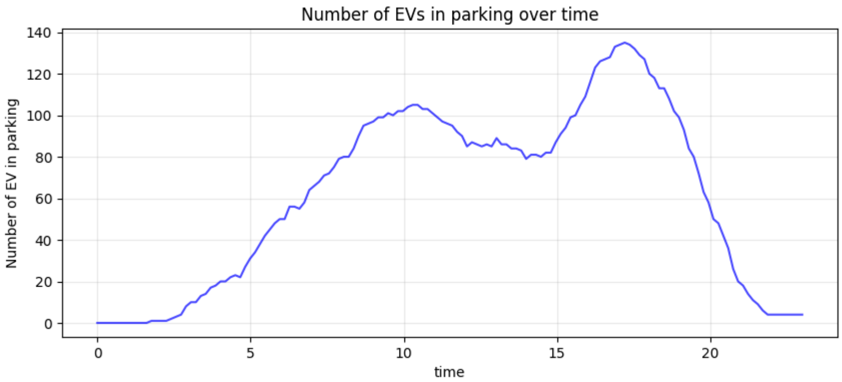

Figure 4.7: Number of EVs in parking over time

The number of EVs in the parking lot varies during the day, as seen in Figure 4.7. The pattern matches with a two-peak arrival model, with a first large wave of arrivals occurring in the morning as people arrive for work, followed by a second, smaller peak in the late afternoon when more cars arrive. The curve then gradually drops as cars begin to leave, and by the end of the day, the lot is nearly empty once more. Because it directly influences the charging demand, this daily routine is significant. The system must manage a high volume of cars arriving quickly during peak hours, which impacts the way charging is allocated. There is greater flexibility to move loads into periods with lower prices or to balance out grid stress during the slower hours. Since charging demand is never distributed evenly over the day, it is essential to know about these availability fluctuations in order to test how well the model fits reality.

During the day, Figure 4.8 displays the dynamic behavior of CO_2 and electricity prices. The two curves show common daily patterns of consumption, with varying heights at midday and late afternoon. Electricity costs rise during times of higher demand, while CO_2 prices vary according to the grid's generation mix, increasing when fossil fuels are used most and decreasing when renewables are easier to access.

The optimizer makes charging decisions based on these profiles. The model can postpone or move charging to more cost-effective and environmentally friendly hours by identifying when expenses or emissions increase. In actuality, this includes being sure that cars get the energy they require while avoiding times when emissions are high or electricity costs are high. The model is able to reach a balance between environmental responsibility and cost savings because of this dynamic adaptation.

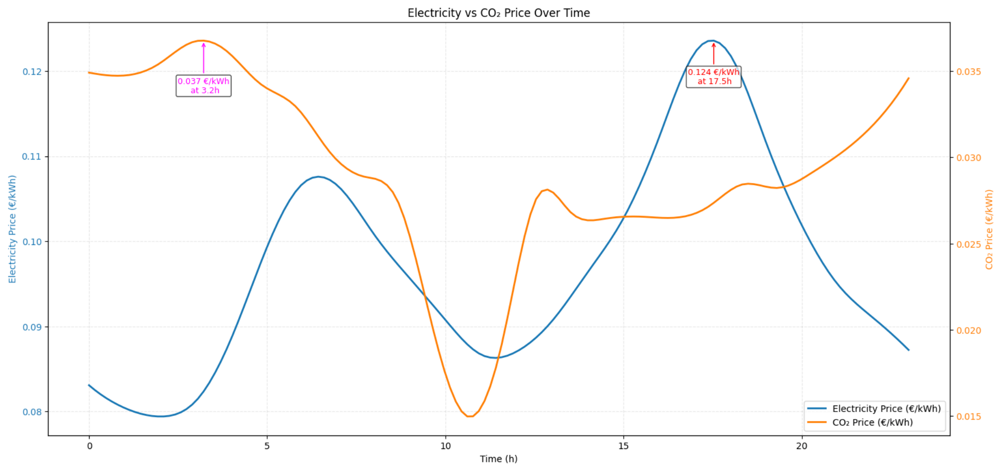

Figure 4.8: Electricity and CO_2 Prices Over Time

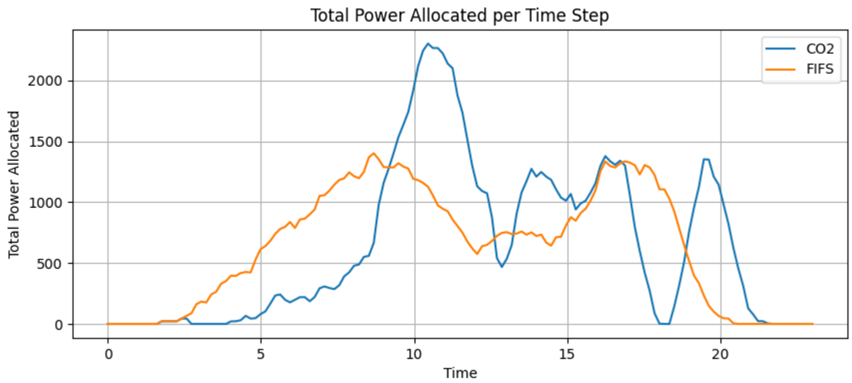

Figure 4.9: Total Power Allocation – CO_2 Optimal vs. FIFS

Power allocation between the baseline First In First Served (FIFS) method and the optimal emissions aware model is compared in Figure 4.9. The most efficient plan option avoids early morning peaks and delays charging until times when emissions and costs are lower. The FIFS approach, on the other hand, ignores emission trends and begins charging as soon as the user arrives.

For evaluating the trade off between charging cost and CO_2 emission changes,Figure 4.10 illustrates how these factors change as the trade off coefficient λ increases. When λ is small, the model focuses mainly on keeping costs low, which results in higher emissions. As λ grows, the model starts to prioritize reducing emissions, and we can clearly see CO_2 output drop. However, this comes at the expense of higher charging costs.

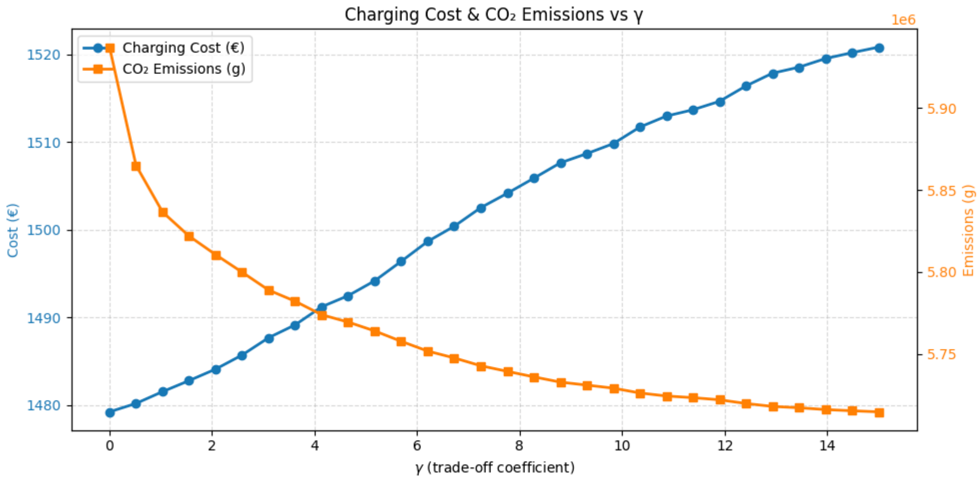

Figure 4.10: Trade-off between Charging Cost and CO_2 Emissions as λ varies

Another significant meaning is shown by the curves' shape. Giving emissions more weight initially results in significant reductions with only a small cost increase. However, the emission reductions decrease as λ increases, and the additional cost continues to rise. In other words, there are decreasing benefits, meaning that reducing emissions further becomes significantly costly after a certain point. It emphasizes the importance of balance. Decision makers can select the λ value that best fits their objectives based on whether reducing emissions or saving money is the top priority.

However, when we find a balance between cost and environmental impact, the CO_2 Model shows how EV charging can be handled more effectively. In contrast to the common First In First Served (FIFS) approach, which only charges cars as they arrive, this model takes into account external influences such as emission levels and electricity prices. It is therefore a good fit for strategies that are focused on sustainability because it more efficiently schedules charging, lowering costs and emissions. Cost and emissions are combined into a single goal in the emissions aware case. The model looks for a compromise that takes into account rather than treating them independently.Here our outcomes shows how well this strategy works. The optimized model's total charging cost was 1508.059 €, while the FIFS strategy's cost was 1569.445 €. This highlights the Emission CO_2 Model's strength. It indicates that when optimization is used, charging can be both cost-effective and environmentally beneficial. Techniques like these will become more crucial as EV adoption rises in order for building a charging system that is both economical and in line with lowering carbon emissions

4.3 Cohort Model Results

Simulation Setup

The Cohort Model is analyzed over a 24-hour period, split into 144 time steps of 10 minutes each. Here our goal is instead of following each EV individually, the vehicles are grouped into $n = 10$ energy classes based on their remaining charge, ranging from class 0 (fully charged) to class 10 (empty). The model greatly simplifies the management of large fleets of EVs by making charging decisions for groups rather than individuals.

In order to simulate a standard charging facility with reliable infrastructure, each charging point delivers a constant power of $P_0 = 22$ kW, and the total power available at each time step is fixed at $\bar{P}(t) = 3300$ kW.

Figure 4.11 shows that vehicle arrivals follow a realistic, time dependent Poisson process $\lambda(t)$, which indicates two main daily charging peaks: one in the evening (around 5 PM) for residential use and one in the morning (around 6 AM) for workplace charging. A Poisson distribution with mean $\lambda(t)$ is used to determine the number of new EVs at each time step. considering that no EV arrives fully charged, these arrivals are then allocated in fixed, increasing amounts among the non-zero demand classes.

The probability that a vehicle is assigned to class i is given by:

$$z_0 = 0, \quad z_i = \frac{i+1}{\sum_{j=1}^{n}(j+1)} \quad \text{for } i = 1, \dots, n$$

Then, the number of arrivals in class i at time t is computed as:

$$a_i(t) = \lfloor z_i \cdot \Delta \cdot \psi(t) + 0.5 \rfloor, \quad \psi(t) \sim \text{Poisson}(\lambda(t))$$

Departures are modeled with a straightforward rule, vehicles that are closer to being fully charged (in the lower demand classes) have a higher chance of leaving. The probability that a vehicle in class i departs is given by:

$$\alpha_i = \frac{n+1-i}{n+2}, \quad \text{for } i = 0, \dots, n$$

Historical Italian market data (January 2023–March 2024) is used to generate the electricity price vector π_t. A multivariate Gaussian distribution is used to model the price in order to simplify stochastic simulations in the future. But we use its average profile for this experiment.

We use a dissatisfaction vector $\beta = (0, 1, \dots, 1)$, meaning that mostly all classes except the fully charged class between cost minimization and fair charging. By setting the weight of dissatisfaction to $\gamma = 1.3$, the system prevents from neglecting users with high energy demands while still trying to reduce the overall cost.

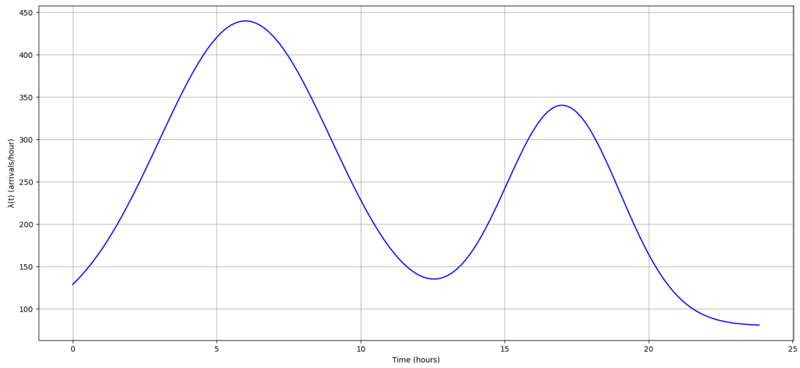

Figure 4.11: Time-varying arrival rate $\lambda(t)$ of EVs (arrivals/hour)

EVs' Poisson arrival rate is shown in Figure 4.11. The curve identifies two main activity waves. The first is an important increase in the morning, which is caused by cars arriving at public charging stations or places to work early in the day. Early evening is when the second, even more obvious peak happens, signifying an increase in residential or after-work charging as people return home.

These dynamics indicate what is frequently seen in actual urban charging environments, so they are not just theoretical. Demand fluctuates over the day based on traffic patterns, and households plugging in their electric vehicles (EVs) simultaneously with other high consumption appliances in the evening creates an additional challenge to the grid. It is crucial to include this behavior in the model because it enables us to evaluate how different charging methods function in real-world scenarios where demand is fluctuating rather than distributed equally throughout the day.

Results and Discussion

The testing environment for the cohort based smart charging model tries to model real world market signals and realistic vehicle behavior. This means that the model is not static but rather responds to outside variables such as the arrival of EVs during the day and the changes in electricity prices over time. To determine how well the optimizer can adapt to these shifting inputs, it needs to find a balance between efficiency (ensuring that overall energy and cost use remain optimal) and fairness (ensuring that no group of vehicles get unreasonable choice). These dynamics can be seen in the following figures, which show how the model modifies charging choices based on user behavior and grid signals.

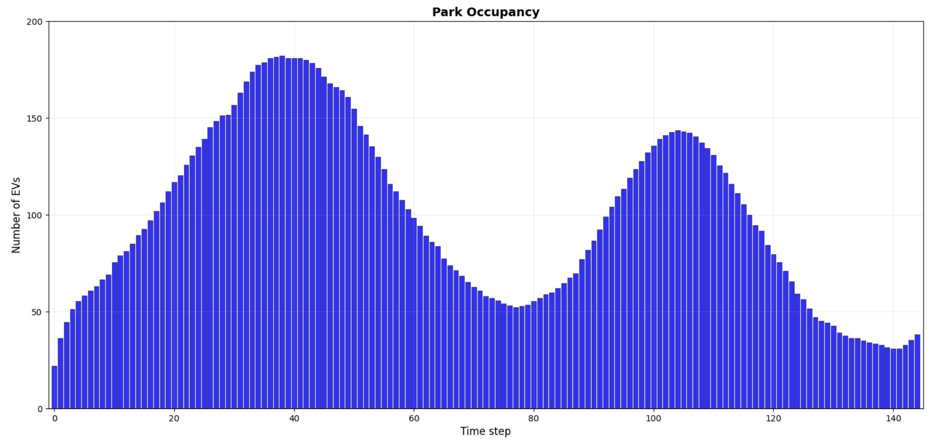

Figure 4.12: Total number of EVs parked over time.

We can observe the total number of EVs present in the system at each time step in Figure 4.12 . The parking occupancy naturally follows the arrival profile, with peaks in the morning and evening and a dip in the afternoon. This occupancy trend directly impacts the demand pressure on the charging infrastructure.

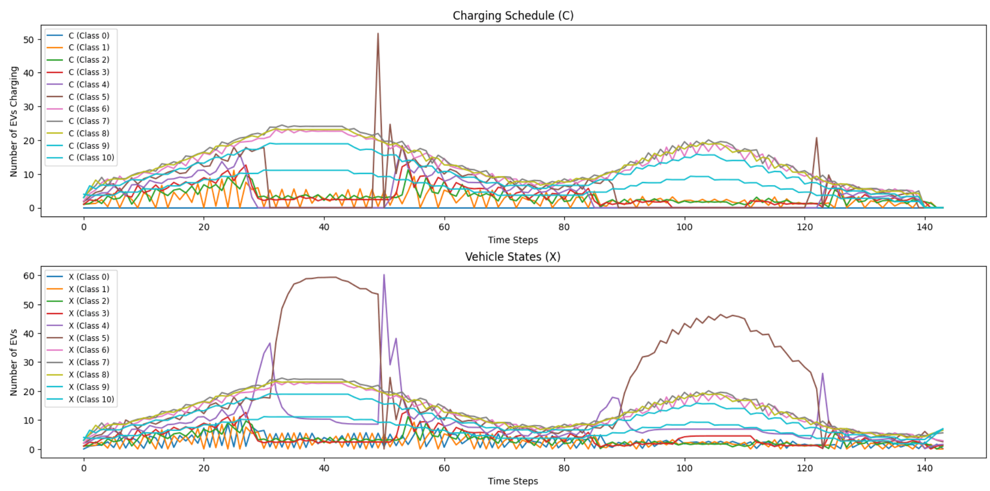

Figure 4.13: Charging schedule (top: $C_i(t)$) and vehicle states (bottom: $X_i(t)$) across classes.

A detailed look at the cohort level dynamics is provided in Figure 4.13 :

- The distribution of charging across demand classes is shown in the top plot. The attention of our model is the efficiency that can be observed in a more reasonable charging of higher classes (vehicles with higher energy needs), especially during times of low prices.

- The evolution of vehicle states appear in the bottom plot. As time goes on, EVs transition from classes with higher demand to those with lower demand, showing how energy is delivered gradually and how cars depart after being fully charged.

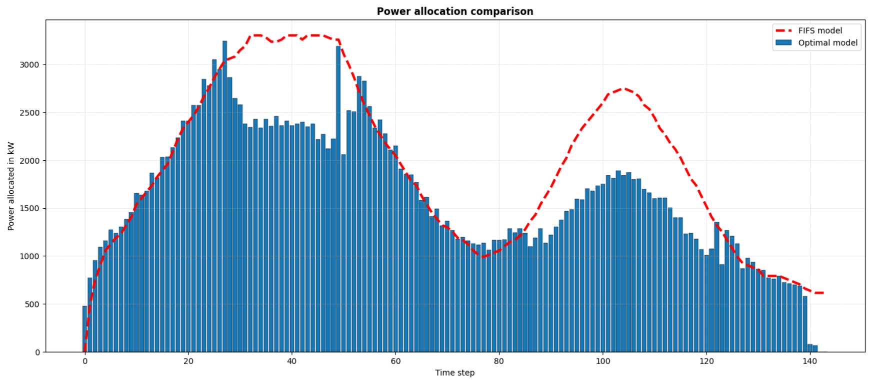

Figure 4.14: Power allocation comparison between the optimal cohort model and the FIFS baseline.

Finally,Figure 4.14 compared the cohort optimizer's overall power allocation with the more basic First-In-First-Served (FIFS) method. Vehicles are charged immediately after arrival in FIFS, regardless of grid conditions or electricity prices. In contrast, the cohort model avoids costly price peaks by moving charging timing at a time, therefore taking better choices. This leads to a much smoother and more cost effective power profile, one that aligns better with the valleys in the electricity price curve.

The numbers clearly show the benefit of this strategy, the optimal cohort solution reached a total cost of 3856.108 €, whereas the baseline FIFS method ended up at 4611.79 €. This difference highlights just how effective class based approach can be in smart charging. By responding to both price signals and class dynamics, the Cohort Model proves it can scale to large groups of EVs while keeping the system efficient.

4.4 Comparison and Discussion

This section summarizes and compares the performance of the three optimization models *Power Allocation*, *Emission CO_2*, and *Cohort* against the *First In First Served (FIFS)* baseline, based on the results observed in the simulation studies.

The FIFS is a naïve baseline that is not sensitive to system load, price, or emissions. EVs are just charged in the order that they arrive, which results in wasteful power consumption and increased expenses,

The **Power Allocation Model** effectively reduces electricity expenses by smartly moving charging to off peak hours. Real cost savings can be achieved by its dynamic response to changing electricity prices. But because of its clear trade-off between cost and user satisfaction, this model sometimes leaves a small percentage of EVs partially undercharged, particularly during price peaks. A dissatisfaction penalty γ is introduced to improve fairness.

By integrating environmental factors, the **Emission $\mathbf{CO_2}$ Model** expands on the cost-based formulation. By optimizing for cost and emissions, it decreases charging activity even when prices are low during times of high CO_2 intensity. The environmental benefit exceeds the trade-off, even though this may occasionally result in somewhat higher energy costs than the only cost-minimizing model. However, the model objective benefit showed in our simulation, as it improved over FIFS in both cost and CO_2 awareness.

Large scale implementations are especially well-suited for the **Cohort Model**. Scheduling is made easier and scalability is improved by classifying EVs according to demand. It makes sure that vehicles with greater demand are given priority earlier in the day and effectively handles class changes.

Summary of Insights:

- *Power Allocation Model*: Best for cost driven settings with moderate computational effort. Sensitive to γ tuning.
- *Emission CO_2 Model*: Ideal when sustainability goals are prioritized alongside cost.
- *Cohort Model*: Most scalable and efficient for large deployments; delivers strong performance under uncertainty and population level dynamics.

These results illustrate that optimization, when properly formulated, can provide substantial benefits in cost savings, emission reduction, and operational fairness. The choice of model depends on the scale of the system, computational resources, and policy goals.

Chapter 5

Digital Twin panel

5.1 Introduction

To maintain sustainable and efficient operations, real time monitoring and future behavior prediction tools must be developed due to the complexity of EV charging systems. We created a digital version of the smart charging models described in Chapter 3 in order to satisfy this demand. This digital twin serves as a simulation of the real system, allowing users to change parameters and see how the models respond in different situations.

The capabilities, architecture, and implementation of the Python based digital twin system which includes a specially created graphical control panel for real-time scenario exploration—are presented in this chapter.

5.2 System Architecture

The digital twin is structured into modular components:

- **Model Core:** Integrates the three optimization models: Power Allocation, Emission CO_2, and Cohort with an interface.
- **Parameter Input Engine:** Allows users to select values such as number of EVs, time resolution, price type, and penalty coefficients.
- **Scenario Engine:** Supports both historical data and stochastic price profile via Monte Carlo simulations.
- **Visualization Panel:** Displays charging behavior, power allocation, occupancy, cost, and emissions over time using Matplotlib and Customtkinter.

Design Choices and Rationale

One of the main goals in developing the digital twin was to make it both user-friendly and useful for research. For the graphical interface, we decided to use the Python `CustomTkinter` library. We considered alternatives such as Dash or Streamlit, but those are more web-focused and would have required setting up a server. Since our priority was to keep the tool lightweight and easy to run locally, `CustomTkinter` turned out to be the best option. It also works smoothly with Matplotlib, which we were already using for visualizations, allowing the whole system to be managed within a single Python workflow.

The system's modular design makes it much easier to grow and maintain over time. To enable separate development and testing, the framework's Model Core, Input Engine, Scenario Engine, and Visualization Panel are all kept separate. For instance, updating the scenario generator does not require us to touch the visualization panel. During development, it was very helpful that we could debug and refine individual modules before integrating them into a complete system.

How the Components Interact

There is a step-by-step description of the workflow between the modules. Because of the way the system is set up, every model and its logic is coded in a different function. Through the control panel created with `CustomTkinter`, a main script then connects everything. Following the user's selection of a model and plots, the script executes the relevant function and presents the visualizations and outputs.

The user can choose at the outset whether to use scenario based profiles or historical data for price inputs. The script calls the price generation function first, chooses the relevant scenario, and then uses that data to run the model if the scenario option is selected. The workflow remains straightforward but adaptable in this way: models can be executed separately, and the panel offers a simple means of interacting with the system and viewing results right away.

Why a Digital Twin Panel and Not Just Scripts?

We now think about why developing a digital twin panel platform was necessary rather than just employing Python scripts to run the optimization models. Usability is the main cause. Scripts are useful for researchers with coding experience, but they are not very useful for operators, policymakers, or students who just want to test scenarios without working with the code. This issue is resolved by the digital twin, which offers an interactive environment in which users immediately see the results in plots and change parameters as needed. This bridges the gap between theoretical models and practical applications and simplifies significantly optimization concepts.

5.3 Interactive Control Panel

Using Python's CustomTkinter library, we created a unique graphical user interface to make the digital twin system easier to use. With the GUI, users can:

- Select the optimization model to simulate.
- Choose between historical or scenario based electricity prices.
- Modify simulation parameters such as:
  - Time steps (Δ)
  - Number of EVs (N)
  - Penalty weight (γ)
  - Number of classes for the cohort model
  - Grid and Socket Capacity
- Visualize the results: power allocation, number of EVs, cost per EV, emissions, and class transitions.

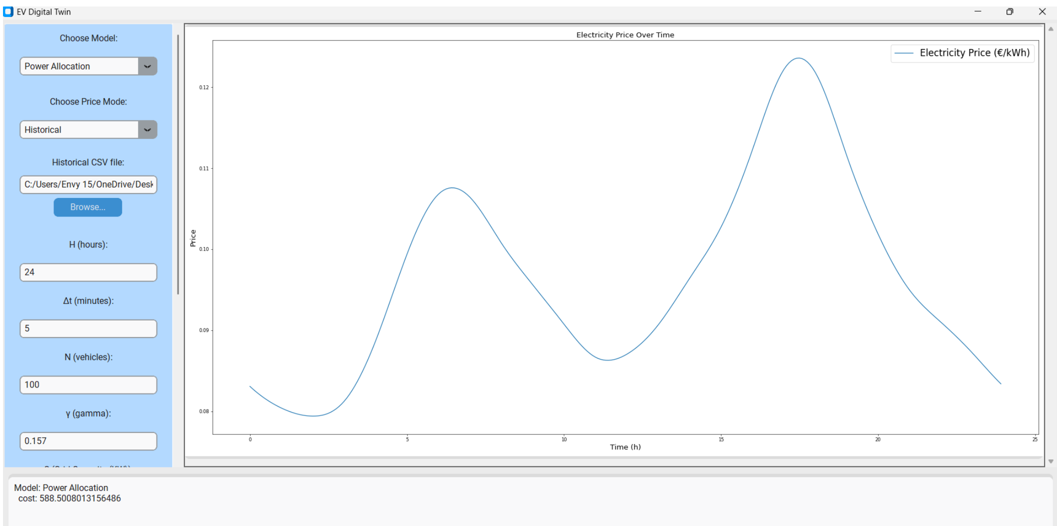

Figure 5.1: Python GUI control panel of the EV Digital Twin input parameters (left) and real time simulation result (right).

The GUI created for the EV Digital Twin is shown in Figure 5.1. Users can modify the model type, time step, fleet size, grid capacity, and penalty factors in the left panel. The dynamic output, such as the electricity price curve over the simulation horizon, is displayed on the right panel. This configuration facilitates scenario testing and allows for the visualization of model behavior in various scenarios.

The results panel automatically updates when the parameters are changed and the simulation is rerun, and the interface has a modular design.

Output Visualizations and Feedback

The developed Digital Twin's ability to visually represent each smart charging model's behavior under different input conditions is one of its main advantages. Through dynamic plots, the system offers real time visual feedback after simulation parameters are set, and a model and pricing mode are chosen.

Depending on its goal, each model generates different results.

Power Allocation

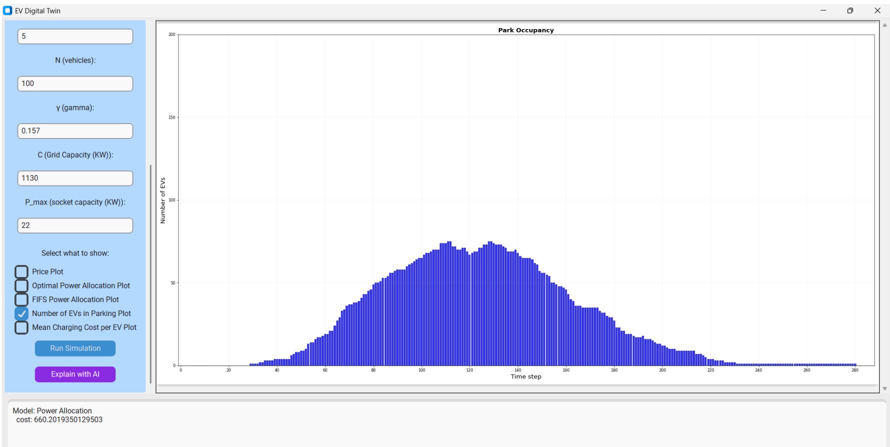

Figure 5.2: EV parking occupancy pattern in Power Allocation model scenario.

The interface in the **Power Allocation Model** displays how charging loads are moved to less expensive hours in order to reduce electricity expenses. The optimized power profile effectively avoids price peaks, as shown in Figure 5.1, while the more basic First-In-First-Served (FIFS) approach is unable to adjust and results in a higher total cost. Additionally, Figure 5.2 shows how parking occupancy varies over time, highlight times of high demand and the way the system responds to cars coming and going. This view is particularly useful for testing how charging behavior is affected by different γ values or time resolutions.

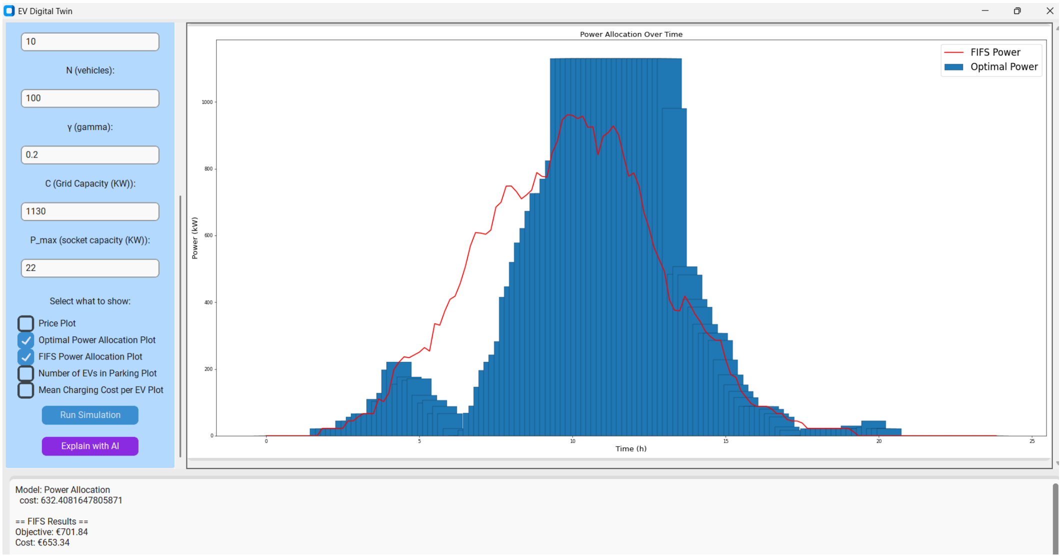

Figure 5.3: Power Allocation model with $\Delta t = 10$ minutes and $\gamma = 0.2$. The optimizer strongly prioritizes cost minimization, leaving some user demand unmet.

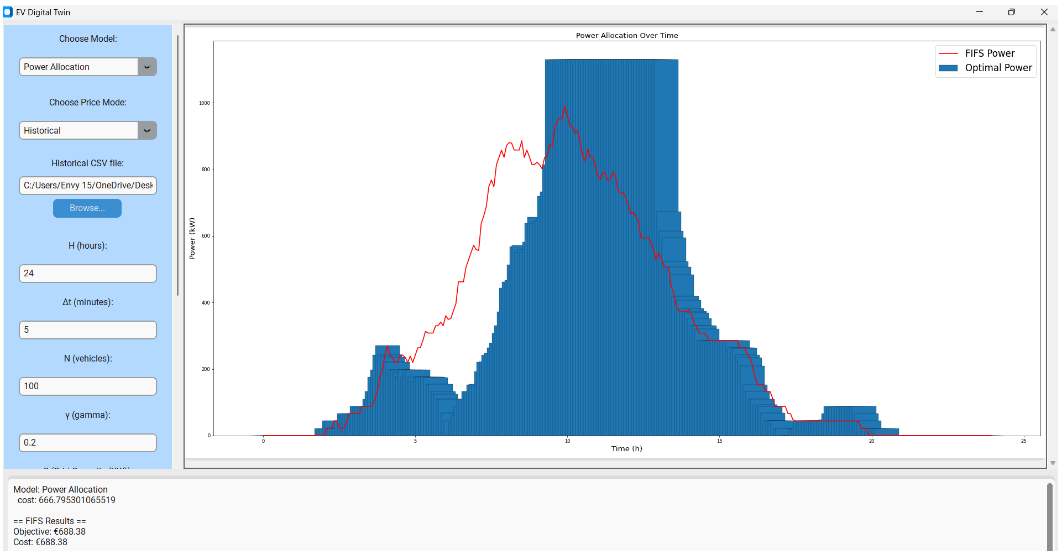

Figure 5.4: Power Allocation model with $\Delta t = 5$ minutes and $\gamma = 0.2$. Finer resolution provides smoother allocation decisions, though cost remains the main driver.

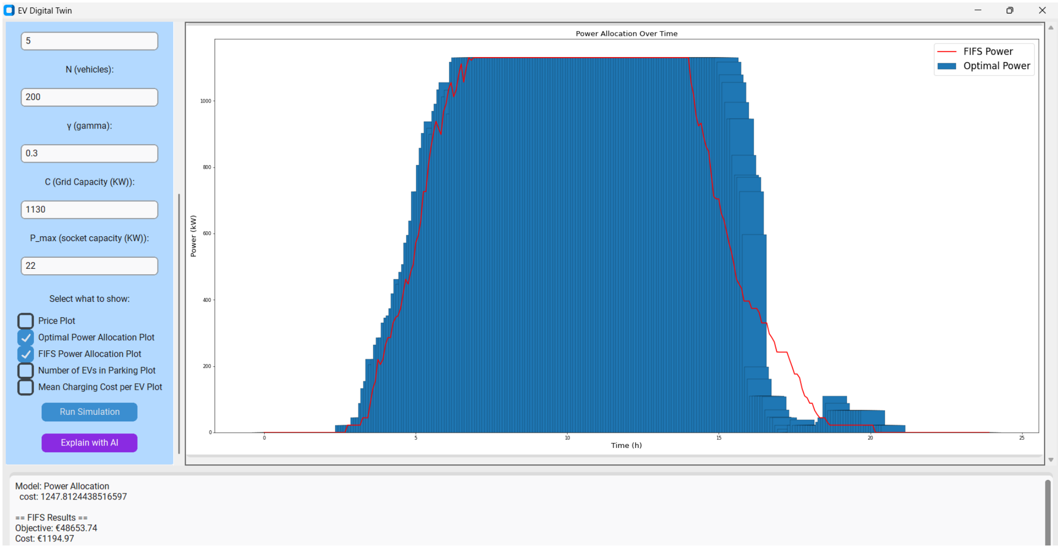

Figure 5.5: Power Allocation model with $N = 200$ and $\gamma = 0.3$. A slightly higher dissatisfaction weight shifts energy earlier to improve user comfort.

The Power Allocation Model's response to changes in the dissatisfaction weight γ and the time resolution Δt is shown in these figures.

The optimizer mainly focuses on cost reduction when γ is set low (Figures 5.3 and 5.4). This results in a significant shift in charging toward the most economical hours, but not all user demands are satisfied. Here, the trade-off becomes clear: while the system maintains low costs, service quality is negatively impacted.

We can compare $\Delta t = 10$ minutes (Figure 5.3) with $\Delta t = 5$ minutes (Figure 5.4) to examine the impact of time resolution. The allocation profile appears smoother and more accurately reflects charging behavior with smaller time steps. In practice, this means shorter time intervals give the system more flexibility to follow price changes closely, though this comes at the cost of higher computational effort.

The model responds by scheduling more charging earlier in the day to decrease dissatisfaction when both the fleet size and γ are increased (Figure 5.5). The effect of $\gamma = 0.3$ is evident: the optimizer balances price and user comfort by allocating a certain amount of the demand earlier rather than waiting only for low priced hours. Small changes in γ have a greater impact due to the larger number of electric vehicles ($N = 200$), which also increases overall effects.

The optimizer better distributes charging when the grid capacity is limited to 1000 kW (Figure 5.6), but it is unable to completely match demand peaks. Consequently, in comparison to situations with greater capacity, certain vehicles suffer from greater deficits.

The model becomes more flexible when the grid capacity is increased to 1500 kW (Figure 5.7). In order to minimize deficits and enhance service quality, the

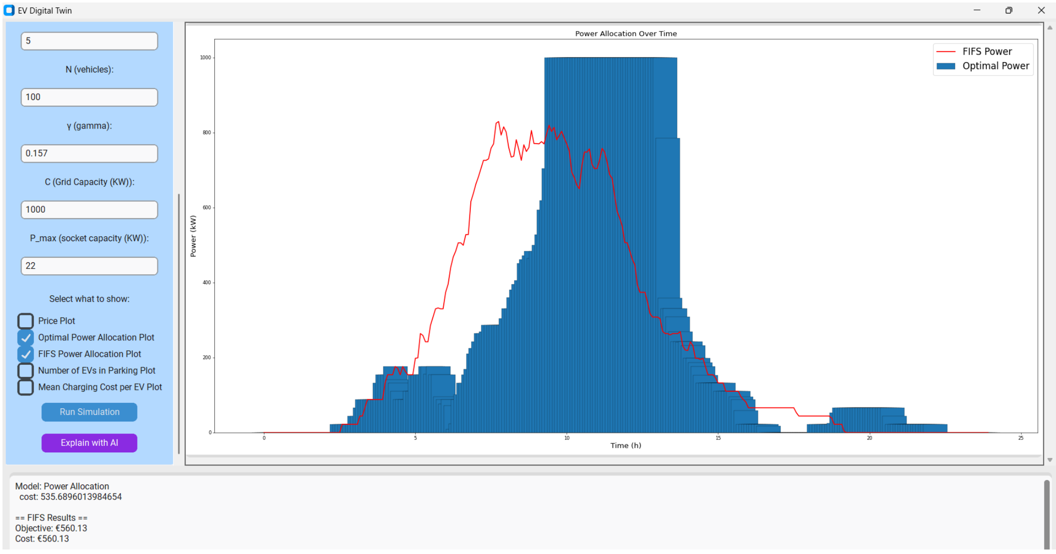

Figure 5.6: Power Allocation model with grid capacity set to 1000 kW.

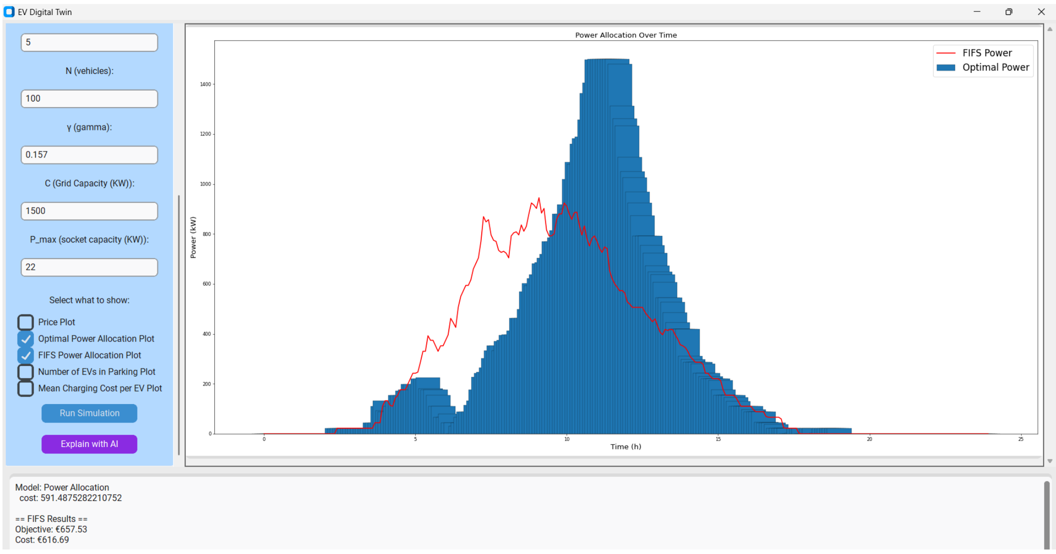

Figure 5.7: Power Allocation model with grid capacity set to 1500 kW.

allocation carefully tracks vehicle arrivals and electricity prices.

All things considered, these results indicate how heavily the parameters selected affect the Power Allocation Model's behavior. The balance between cost reduction and user satisfaction is determined by the dissatisfaction weight γ; low values motivate the optimizer to reduce costs, while higher values encourage earlier charging to prevent demand from remaining ignored. The time step Δt is also important because, although it requires more computation, smaller intervals enable more precise and seamless charging adjustments. Because the effects of allocation strategies are increased at all levels in larger systems, fleet size makes these trade-offs even more clear.

Grid capacity is an additional crucial factor. The model is forced to schedule more precisely when capacity is limited, and some demand might go unused. Higher capacity allows charging to be distributed more uniformly, remaining closer to user demands and price signals. When combined, these experiments demonstrate how the Power Allocation Model's parameters have a significant impact on how charging strategies are adjusted to actual conditions.

Emission CO_2

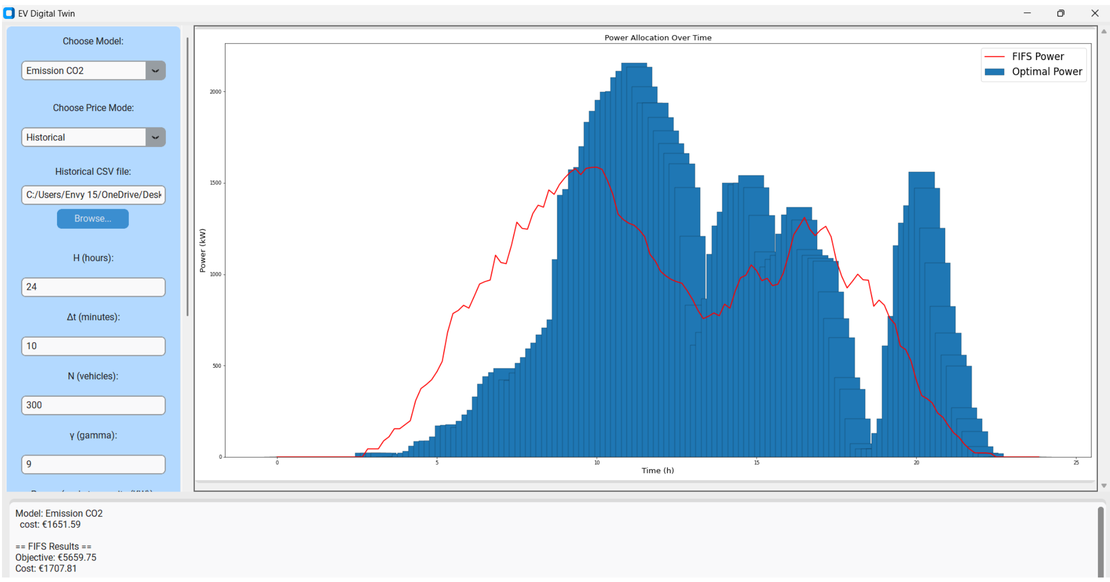

Figure 5.8: Emission CO_2 model: optimal vs. FIFS strategy.

For the **Emission CO_2 Model**, the control panel shows a clear comparison between the optimized charging profile and the baseline FIFS strategy. As shown in Figure 5.8, the optimized approach concentrates power usage during cleaner grid periods, reducing both emissions and costs. The blue bars represent the optimized schedule, while the red line shows the less efficient FIFS charging pattern.

Figures 5.9 and 5.10 show the response of the Emission CO_2 model to increasing fleet size and dissatisfaction weight. The optimizer begins shifting charging earlier in the day with $N = 200$ and $\gamma = 5$, ensuring that more vehicles depart with enough charge while remaining true to emission signals. This effect becomes clearer when the fleet reaches $N = 500$ and $\gamma = 10$. This is because the model allocates a

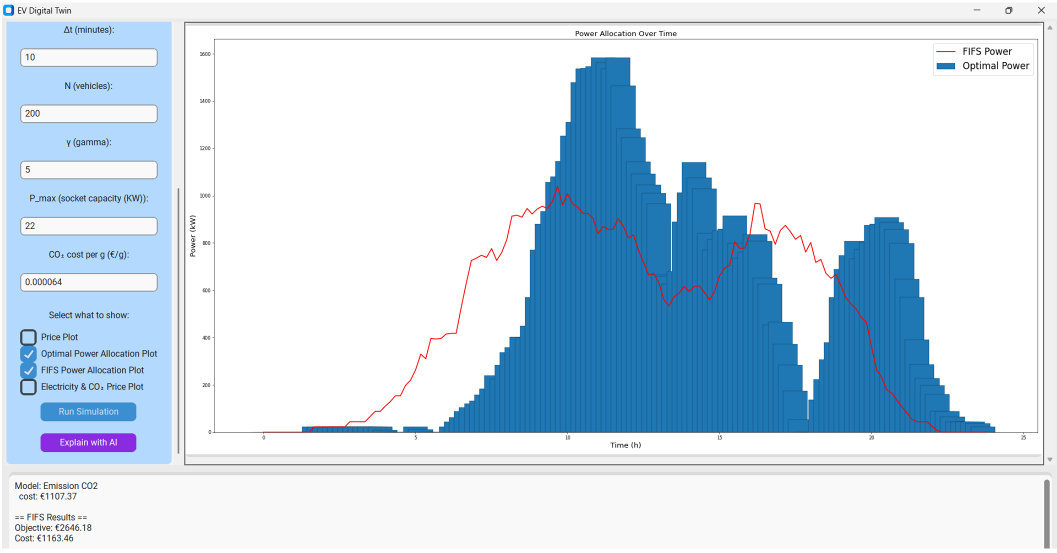

Figure 5.9: Emission CO_2 model with $N = 200$ vehicles and $\gamma = 5$. The optimizer shifts charging away from high-emission hours, balancing cost and sustainability.

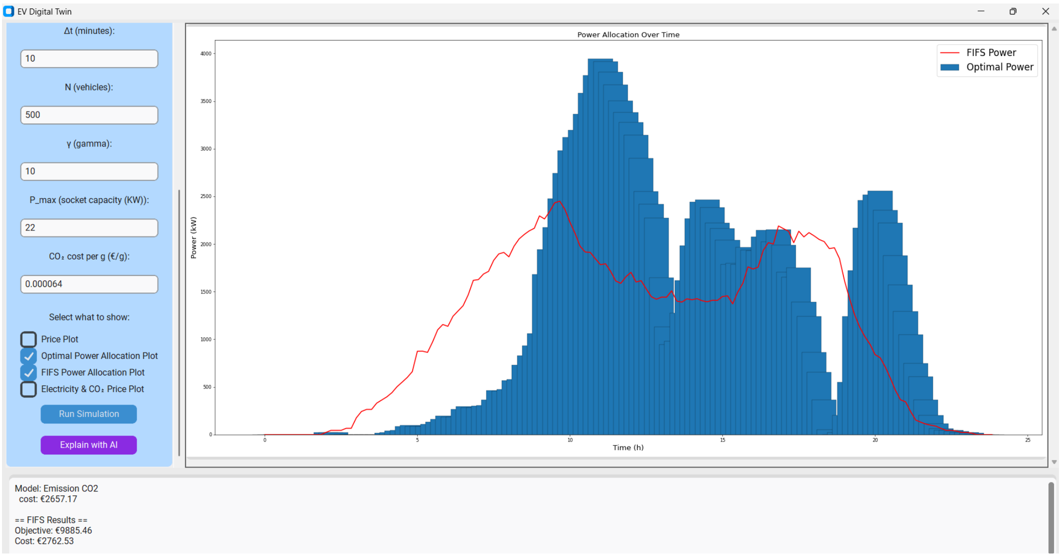

Figure 5.10: Emission CO_2 model with $N = 500$ vehicles and $\gamma = 10$. Scaling up the fleet and dissatisfaction weight highlights the trade-off between user comfort and emission reduction.

significant amount of power earlier and keeps a higher baseline to meet the increased demand. In the meantime, it continues adjusting to cleaner grid periods. These findings show the model's scalability ,environmental impact and user satisfaction.

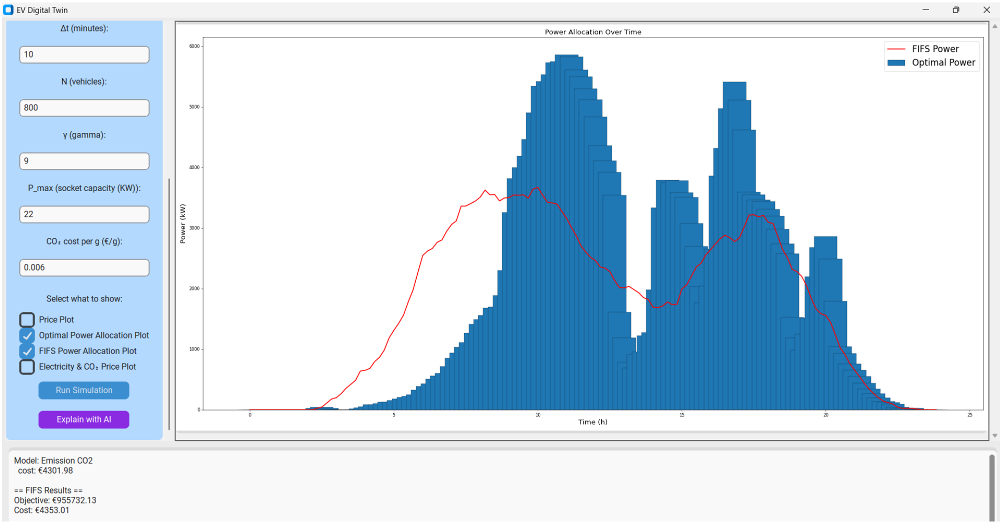

Figure 5.11: Emission CO_2 model with $N = 800$ vehicles, $\gamma = 9$, and CO_2 cost of 0.006 €/g.

In this experiment, a higher dissatisfaction weight ($\gamma = 9$) was combined with an increase in the fleet size ($N = 800$) and the CO_2 cost parameter (0.006 €/g). The optimal model shifts charging aggressively into periods with lower emissions, even if prices are not minimal, proving how environmental objectives dominate when the carbon cost is high. These settings amplify the trade-offs.

The FIFS baseline (red line), in contrast, ignores emission signals and immediately saturates capacity by simply following vehicle arrivals. Costs and emissions increase as a result. The comparison shows how the optimizer prioritizes greener charging schedules when the CO_2 penalty is increased, particularly at scale where unmanaged strategies like FIFS become more inefficient.

Cohort Model

The **Cohort Model** panels show how cars move between charging classes over time. The number of EVs actively charging in each class (C) shows up in the top plot of Figure 5.12, while the total number of EVs per class (X) is indicated in the bottom plot. This dual view provides important insights into fleet dynamics by highlighting the way the system handles departures, allocations, and transitions.

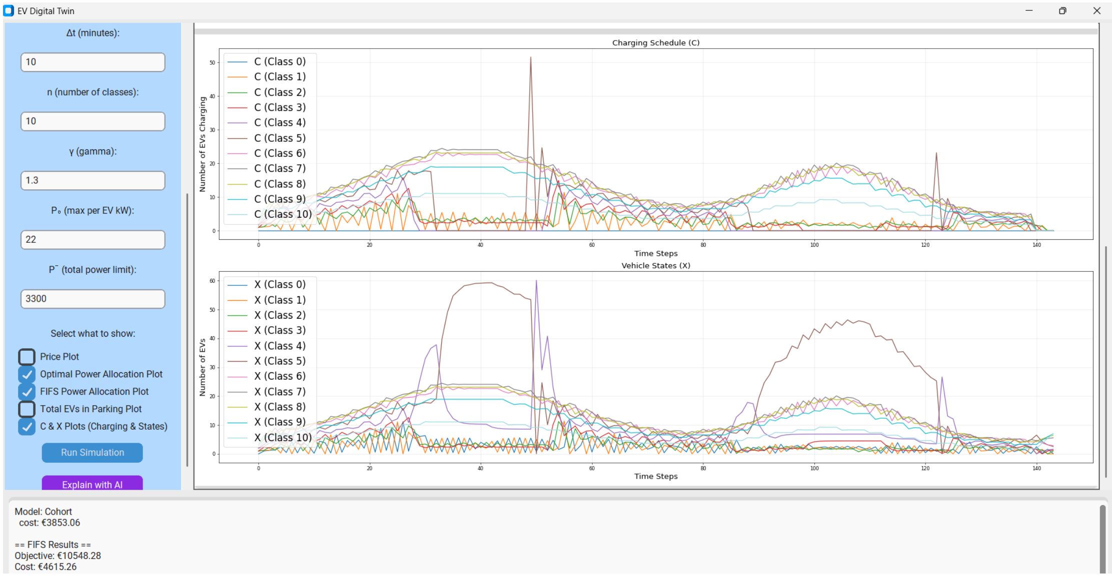

Figure 5.12: Behavior of matrices C and X in the Cohort Model.

We evaluated various degrees of class discretization in order to test the model further. While a larger number ($n = 50$) captures finer detail in energy states and transitions, a smaller number ($n = 10$) groups vehicles more coarsely.

The impact of class resolution is shown in Figures 5.13 and 5.14. The dynamics are easier to calculate and visualize when $n = 10$ is used, but the grouping obscures some of the fleet's variation in energy requirements. Although it requires more computing power, the model gives an accurate picture of how cars transition between classes over time when $n = 50$. Behavior is also influenced by the dissatisfaction parameter γ; higher values encourage the system to charge vehicles earlier, lowering the possibility of unmet demand.

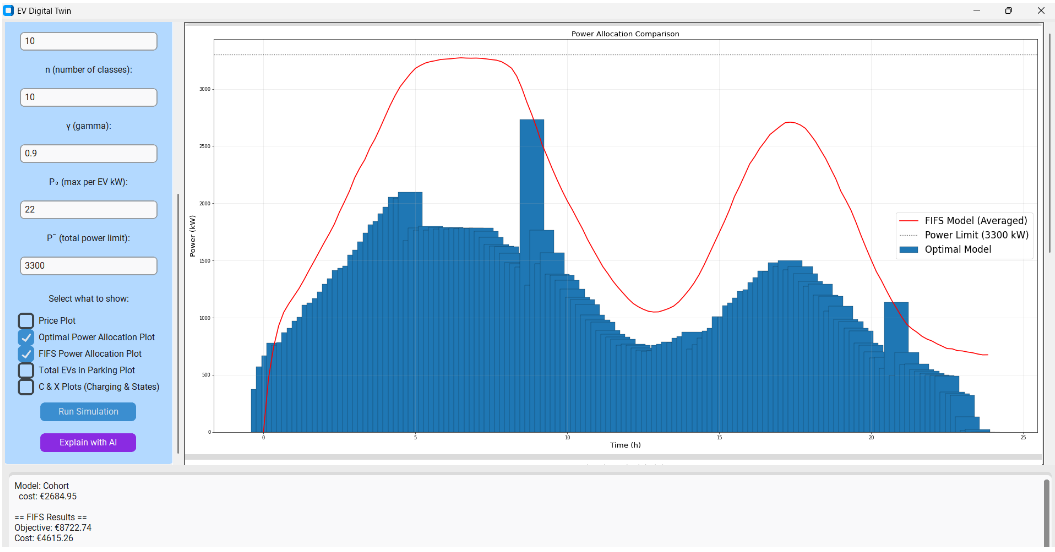

Figure 5.13: Cohort model with $n = 10$ classes and $\gamma = 0.9$. Coarser discretization simplifies the system but reduces resolution of charging dynamics.

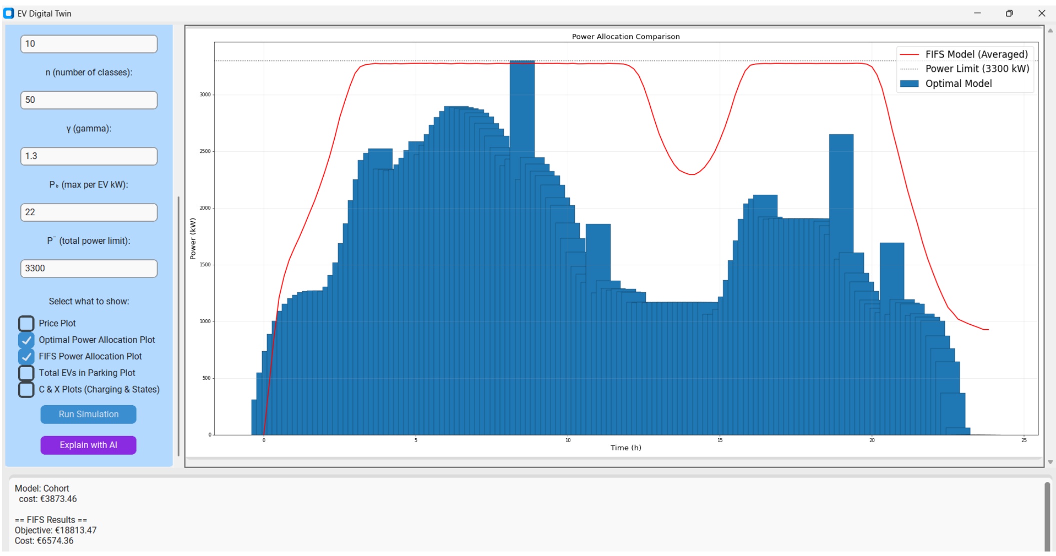

Figure 5.14: Cohort model with $n = 50$ classes and $\gamma = 1.3$. Finer discretization captures smoother transitions and allows more precise scheduling.

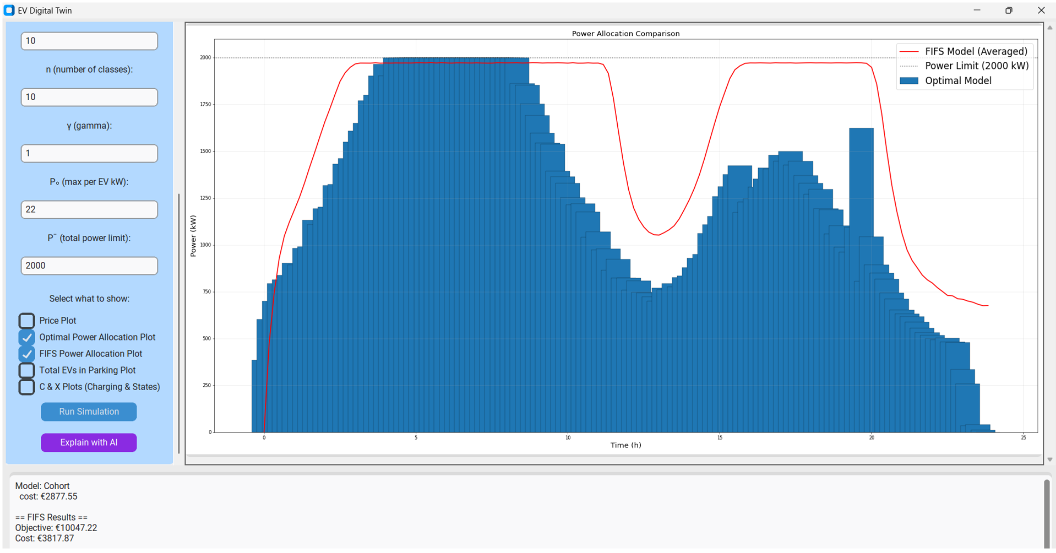

Figure 5.15: Cohort model with $\gamma = 1$ and total power limit $\bar{P} = 2000$ kW. The optimal allocation (blue) adapts to price signals within capacity limits, while FIFS (red) quickly saturates the grid.

Figures 5.15 The cohort model is evaluated with a dissatisfaction weight $\gamma = 1$ and a total grid capacity of 2000 kW. The findings shows how, while staying to the power limit, the ideal approach (blue bars) properly distributes charging over time. The FIFS benchmark, on the other hand, quickly overloads the system, resulting in inefficient allocation and increased expenses.

This experiment shows that the cohort model can manage user demand and cost while managing limited capacity. While the capacity limit emphasizes the system's ability to prevent overloads, something FIFS cannot do, the tunable parameter γ guarantees that charging stays fair across vehicle classes.

Overall, these findings support the flexibility of the cohort approach, which can be adjusted for detail (more classes, smoother dynamics) or simplicity (fewer classes, faster computation), depending on the purpose of the application.

5.4 Scenario Exploration

The ability to perform *what if* experiments is one of the digital twin's greatest benefits. In other words, without putting the real charging system, we can instantly observe how the models respond by adjusting system parameters on the go. Because of this, the twin becomes a secure environment where we can experiment, observe how the system functions, and better comprehend the relationships at play.

You can try a variety of scenarios. For instance, the model begins to prioritize users who want their batteries fully charged before departing if we increase the dissatisfaction parameter γ. This makes it simple to understand how the model achieves a balance between user comfort and other objectives, such as reducing emissions or expenses.

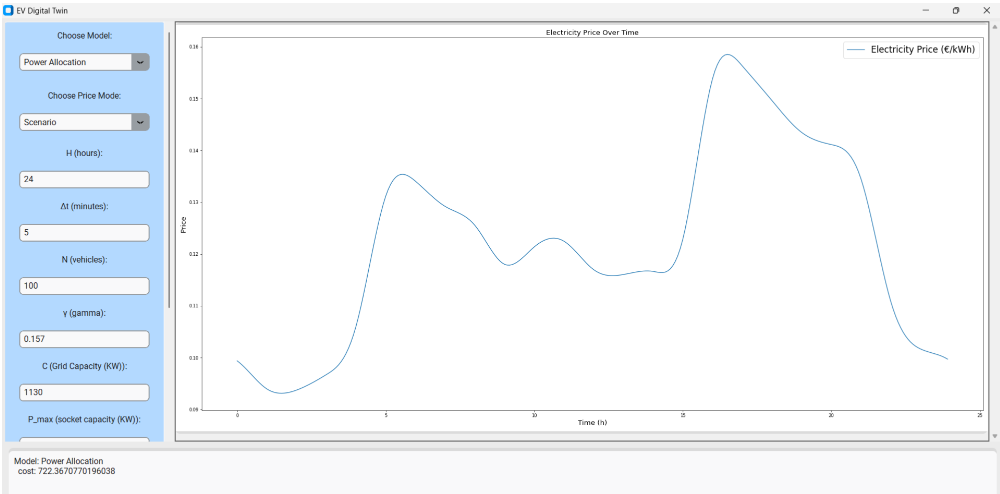

Figure 5.16: Power allocation under a scenario-based electricity price compared with historical data. The model adapts by shifting charging toward cheaper periods.

Changing from historical to scenario-based electricity prices is another scenario. A multivariate Gaussian (normal) distribution is used to generate these prices; the covariance matrix captures the variability and correlation between time periods, while the mean vector represents the historical average. In this way, the digital twin can test how flexible the charging strategies are and simulate realistic changes in market conditions. This situation is shown in Figure 5.16, where charging is moved toward less expensive times while maintaining operational bounds, and power allocation varies based on the price profile.

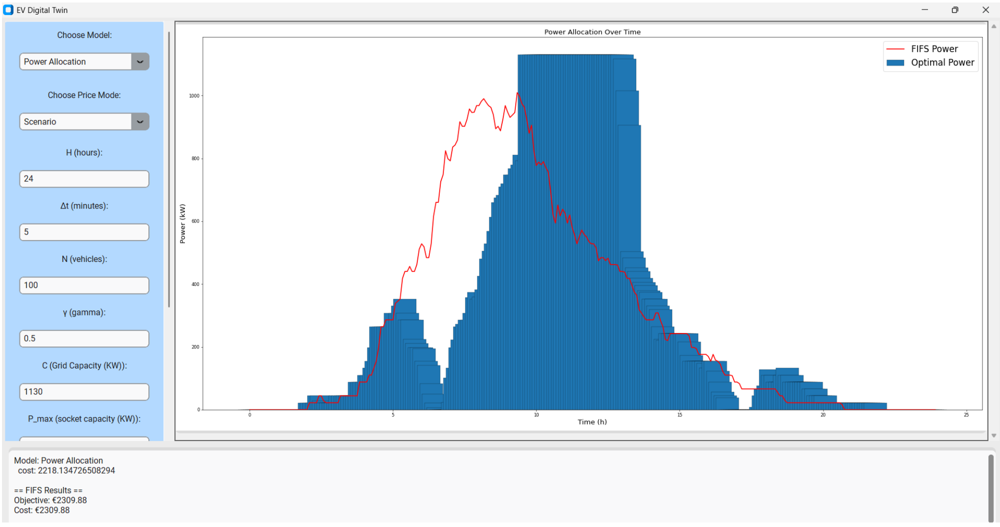

Figure 5.17: Scenario-based Power Allocation with $\gamma = 0.5$. The optimizer balances cost reduction with user satisfaction, scheduling charging earlier compared to pure cost-focused runs.

A different perspective is included in Figure 5.17, where the dissatisfaction weight is set to $\gamma = 0.5$, representing a balanced compromise between ensuring that users depart with sufficient energy and minimizing costs. With this configuration, the optimizer considers driver comfort in addition to price reduction. In order to prevent undercharging of cars with limited parking times, charging is therefore planned to start earlier.

This behavior reacts to the price signals of the scenario-based input while lowering the risk of dissatisfaction in comparison to purely cost driven runs (low γ). Without compromising user confidence in the system, the schedule dynamically adjusts to changes in the market. This balance is crucial for real-world applications because, if the model solely focused on obtaining the lowest prices, users might frequently depart without the energy they require, which could damage trust in smart charging systems. However, the model would ignore price fluctuations and lose a large portion of the economic benefit if it placed too much of emphasis on comfort (very high γ).

This example therefore shows how parameter tuning can create flexible strategies that adapt to both external conditions (electricity markets, emission sign) and internal requirements (user satisfaction, fleet needs).

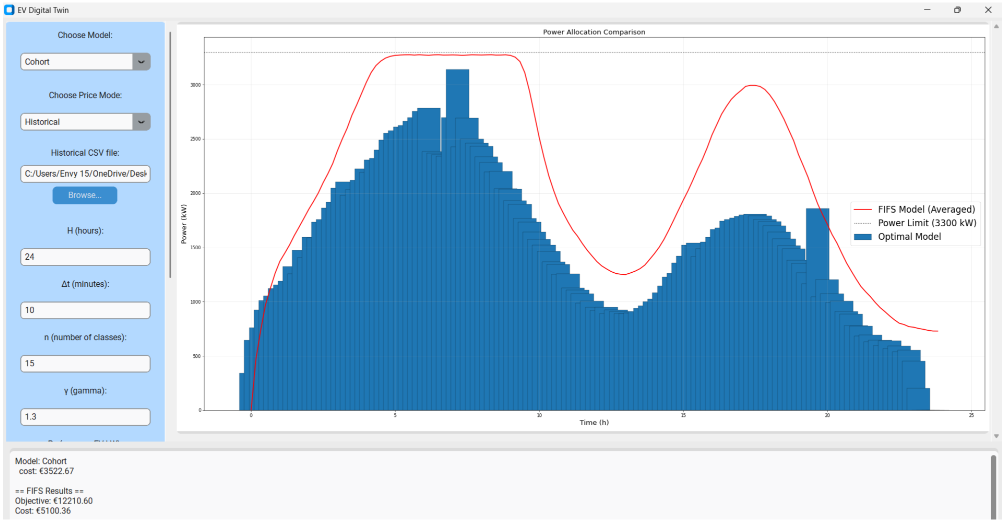

Figure 5.18: Cohort model with $n = 15$ to test the behavior and adaptability of the model.

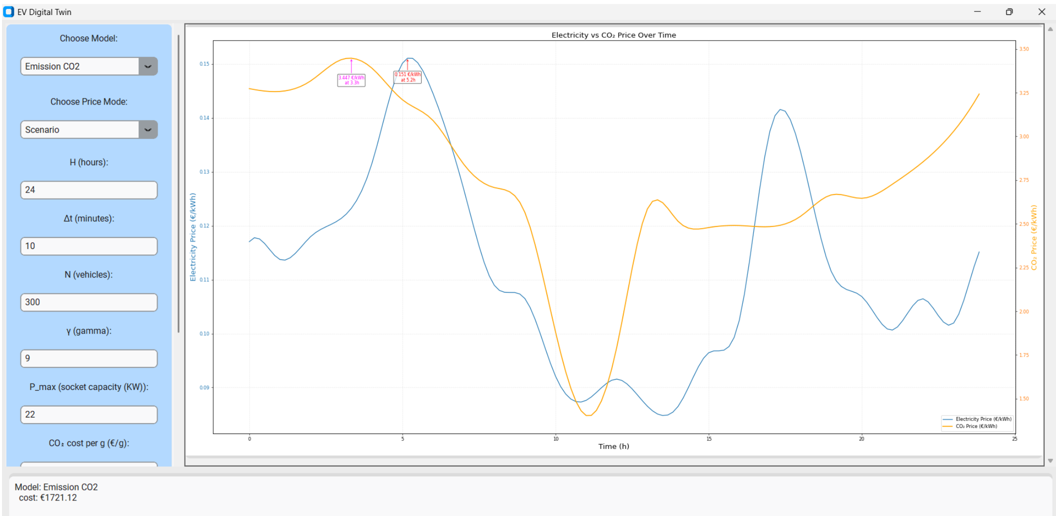

Figure 5.19: Emission CO_2 model under scenario-based prices with CO_2 cost set to 0.006 €/g. Electricity (blue) and CO_2 (orange) prices vary dynamically across the day.

Instead of using historical data, the emission model in this experiment uses a *scenario based* price profile. Realistic fluctuations and correlations are introduced by the multivariate Gaussian distribution used to generate the price of electricity and CO_2. These scenario profiles, as opposed to historical curves, show potential market futures as compared to fixed historical conditions.

Even if electricity is slightly more expensive, the optimizer will strongly prioritize charging during times when the orange line (CO_2 price) is lower because the CO_2 cost parameter is set to a high value (0.006 €/g). As a result, the model becomes more emission-driven and less cost-driven, demonstrating how scenario generation can be used to stress test strategies under future or uncertain price conditions.

The final result indicates how the system adjusts by linking charging with the cleaner windows in the scenario when environmental penalties are high. Unmanaged strategies, such as FIFS, on the other hand, would not respond to these signals, which would result in increased expenses and emissions.

The digital twin is a valuable tool for researchers, system operators, and policymakers due to these types of scenario exploration. Stakeholders can better understand trade-offs and make more informed decisions regarding the management and design of smart charging systems by seeing how the model adjusts to various situations.

5.5 Implementation Overview

A modular and extensible architecture has been employed in the Python development of the digital twin for smart charging of electric vehicles. The system is divided into discrete parts, each of which is in charge of carrying out a particular simulation workflow task. The platform is easier to test, maintain, and add new features to in the future thanks to this design.

CVXPY's ECOS solver, which offers stable results for linear and convex optimization tasks across all three models, is used to solve all optimization problems.

The main script runs a graphical user interface (GUI) created with CustomTkinter and manages the simulation process. This interface allows users to choose a model, set parameters like the time step, the number of vehicles, or the dissatisfaction weight γ, and then see the results in interactive plots. A separate module handles visualization, making sure that results are transparent and comparable across various contexts.

The system also has an experimental AI component that can produce plot explanations. This feature is intended to make it easier for users to understand results, particularly when those results are complex. This is explained in more detail in the following subsection since it is a crucial component of improving the digital twin's usability.

To sum up, the modular architecture guarantees that every part can be created and tested separately while still integrating perfectly into the system as a whole. Because of this, the digital twin can be used as a flexible research tool as well as a useful platform for testing by engineers, operators, and policymakers.

AI Integration

In addition to its optimization and visualization features, the system also includes an experimental AI component, implemented through the OpenAI API. This feature is designed to generate automatic explanations of simulation results. By turning complex graphs into simple narratives, the AI helps bridge the gap between technical outputs and user understanding.

For example, the AI can explain how charging demand peaks align with electricity price fluctuations, or why certain vehicle classes dominate the charging process at specific times. Figure 5.20 shows an example for the CO_2 model, where the AI highlights key patterns in the charging schedule and links them to changes in grid emissions and costs.

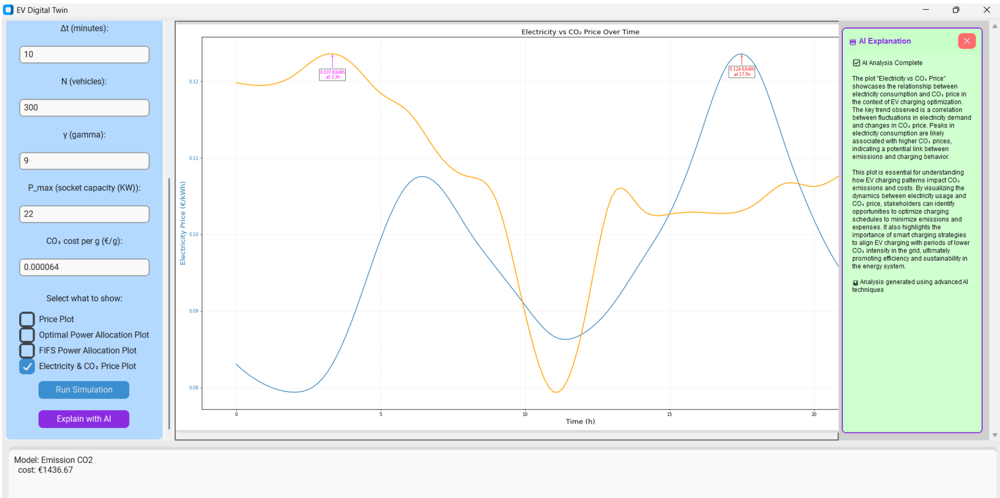

Figure 5.20: AI-generated explanation accompanying a CO_2 price plot. The system provides automatic insights into observed patterns and their meaning for charging optimization.

A second example is shown in Figure 5.21, where the AI interprets the outputs of the cohort model. The top plot (C) shows the number of vehicles actively charging in each class, while the bottom plot (X) shows how many vehicles are in each state overall. The AI highlights important relationships, such as how grid capacity affects simultaneous charging or how high-demand classes dominate the schedule. For system operators, these insights can support better grid management, while fleet managers can use them to improve planning and reduce user dissatisfaction.

To ensure meaningful outputs, the AI assistant only works when a plot is already selected. If a user requests AI feedback without choosing a plot, the system displays an alert message reminding them to first select the relevant visualization. This safeguard prevents empty or misleading explanations and keeps the interaction

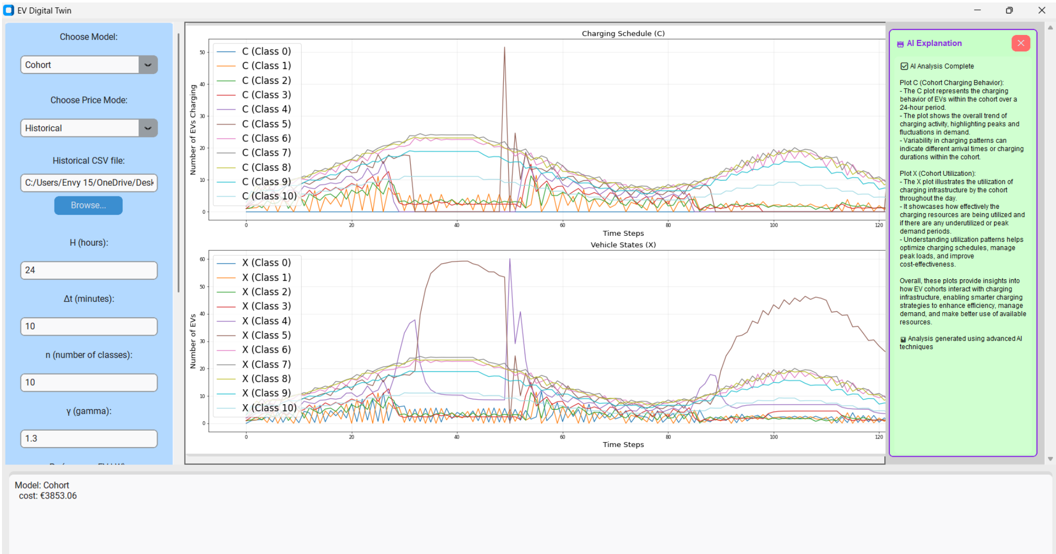

Figure 5.21: AI-generated explanation for the Cohort Model, interpreting the relationship between charging states (X) and active charging classes (C).

intuitive.

Although still in an experimental stage, the AI module shows the potential of combining digital twins with intelligent assistants. It lowers the entry barrier for non technical users, turning the simulator into not only a tool for research and analysis but also a platform for communication and decision support.

5.6 Structure and Usage

The digital twin has been uploaded to a Git repository to make it easier to share and reproduce. The code is organized into clear modules so that each part of the system can be maintained separately, while the graphical interface in `GUI_TK.py` connects everything together into a working application.

File and Folder Overview

These are the primary files and folders:

- `GUI_TK.py`: The system's entry point. It imports every model function and shows the CustomTkinter interface. The user can choose a model, configure parameters, and launch the simulation after running this file. When required, the user can activates the AI explanation module inside the panel.
- `defdeficit_e.py`: Includes the Power Allocation model's implementation.
- `defCO2.py`:The Emission CO_2 optimization model is defined in this file.
- `defCohort.py`: The Cohort model solver is included in this section.
- `defFIFS.py / defFIFS_co2.py`: offer the standard First-In-First-Served tactics for comparison.
- `defA.py, deflam.py, defA_Cohort.py`: Control input generation and arrival procedures for various models.
- `defprice.py`: Creates historical and scenario-based electricity price profiles. The cohort version of the FIFS benchmark is implemented by
- `cohort_fifs.py`: The cohort version of the FIFS benchmark is implemented here.

Executable Packaging with Nuitka

During the development of the digital twin, one of the goals was to create a standalone `.exe` file that could be easily shared with others. This not only makes the tool more user–friendly (no need to install Python or extra packages), but it also prevents direct access to the source code. To achieve this, several adjustments were necessary.

What is Nuitka and Why We Use It

Nuitka is a Python compiler that translates Python code into C++ and then compiles it into native machine code.The result is a standalone executable that runs without having Python installed on the target machine. Compared to interpreted Python, Nuitka offers faster execution and better compatibility with libraries such as `cvxpy`, `scipy`, and `numpy`, which are essential in this project.

What We Achieved

By applying these changes, the final digital twin executable reached the following milestones:

- A distributable `.exe` that works on any Windows machine.
- Stable integration of AI features inside the packaged executable.

With these improvements, the application is now robust, portable, and ready for distribution.

5.7 Summary

The digital twin of this work is not just a simulator. It works as an interactive platform that lets users test out different smart charging methods and observe the outcomes instantly. A clear comparison of cost, environmental impact, and user satisfaction is made possible by the integration of the power allocation, emission CO_2, and cohort based models into a single framework. Users can quickly adjust parameters and see how the system responds with the graphical control panel.

The charts make it clear where optimization works and where traditional strategies, such as First In, First Served, don't work. The tool tests a variety of scenarios, including changes in electricity prices, fleet size adjustments, or stronger comfort rules, to provide researchers, policymakers, and grid operators with accurate, practical guidance on system performance.

One of the main characteristics of the system is the experimental integration of AI. It can generate straightforward plot and trend explanations, making the results easier for non technical users to understand. Along with complex graphs, the system provides bright insights. This feature shows that digital twins can be used to analyze and more easily explain data. Future improvements to the platform could include adding real time data, connecting to charging hardware, or adding more objectives like multi station control and battery health. The foundation for future systems that are feasible and ready for real world deployment may be strengthened by these additions to the digital twin.

Chapter 6

Conclusion

Electric vehicles (EVs) are becoming more common around the world. They bring clear benefits, such as cutting emissions and reducing the use of fossil fuels, but they also create new challenges for energy systems. One of the main challenges is how to design charging strategies that keep costs low, satisfy user needs, and support environmental goals. This thesis worked on this challenge by developing optimization models and building an interactive digital twin system. The digital twin makes it possible to study charging behavior, test different scenarios, and visualize the results in a clear and accessible way.

We started by creating three complementary optimization models, each looking at the problem from a different angle. The first, the *power allocation model*, aimed to reduce charging costs by shifting loads away from expensive periods. The second, the *emission aware model*, added CO_2 intensity into the decision making, showing how charging could be shifted to cleaner grid periods to reduce environmental impact. The third, the *cohort model*, grouped vehicles by energy demand and tracked their transitions over time. This gave better insight into how large fleets can be managed in a scalable way.

All models passed testing under realistic vehicle arrival, time resolution, and power supply conditions after being mathematically described and implemented in Python using CVXPY. The results of numerous tests demonstrated that optimization-based techniques performed significantly better than the simple First In First Served (FIFS) approach. They produced better outcomes in terms of user satisfaction, cost savings, emissions reduction, and change adaptability.

This thesis presented a *digital twin platform* with a graphical control panel to facilitate the use of these models. Users can adjust parameters like the number of electric vehicles or the electricity price mode using this interface, and they can observe the system's response right away. Comparing models and strategies is made simpler with side-by-side visualizations, and scenario testing offers additional flexibility. For instance, we tested robustness using stochastic prices based on

Gaussian distributions, and we demonstrated how the system responds to increased demand by increasing the number of EVs. Because of these characteristics, the simulator becomes more than just a technical tool; it becomes a useful environment for education, experimentation, and decision-making.

Main contributions of this work include:

- Development of three optimization models (cost, emission, cohort) implemented consistently in Python.
- Creation of an interactive and ready-to-use digital twin platform with a CustomTkinter based graphical interface.
- Real time adjustment of parameters and visualization of results such as power allocation, occupancy, cost, and emissions.
- Scenario exploration using both historical and stochastic electricity prices.
- Comparison with the First-In-First-Served baseline, showing the clear advantages of our optimization models.
- Experimental use of AI to automatically explain plots and simulation results.

Another contribution is the AI feature. Despite its early stages, it suggests the potential for collaboration between AI and digital twins. People without a technical background can use the tool more easily because it simplifies complicated plots into brief, understandable explanations. As a result, the simulator serves as a communication tool for a larger audience in addition to being a research tool.

Future research can go in a number of ways. The system would become more useful if it included real-time data from electricity markets and charging stations. Studying urban charging systems would be helpful by expanding the platform to larger networks, such as city-scale simulations. The model might be more realistic if it included more specific elements like vehicle to grid (V2G) features, and battery aging. For actual deployment, decisions involving several charging hubs would also need to be coordinated. Last but not least, enhancing the AI assistant might improve the digital twin's decision support capabilities by helping planners and operators in real time trade off understanding.

In conclusion, this thesis develops the foundation for effective, sustainable, and user focused charging solutions by fusing optimization, a digital twin platform, and an experimental AI assistant. In this way, the work supports the global transition to cleaner and smarter energy systems while also contributing to research.

Bibliography

[1] IEA. *Global EV Outlook 2024*. https://www.iea.org/reports/global-ev-outlook-2024. 2024.

[2] Ellen Kennedy, Hannah Lindsell, and Nick Pesta. «Electric Vehicles Are on the Road to Mass Adoption». In: *Rocky Mountain Institute Insight* (2025). Updated July 29, 2025. URL: https://rmi.org/electric-vehicles-are-on-the-road-to-mass-adoption/.

[3] S. P. M. Deb and A.-S. M. Deb. «Smart Charging: A Comprehensive Review». In: *IEEE Access* (2022).

[4] Andreas Schuller, Bernhard Dietz, Michael Flamme, and Ludwig Karg. «Charging strategies for electric vehicles in smart grids». In: *Energy Procedia* 78 (2015), pp. 207–212. DOI: 10.1016/j.egypro.2015.11.621.

[5] Yijia Cao, Shengwei Tang, Canbing Li, Peng Zhang, Yi Tan, and Zhikun Zhang. «An Optimized EV Charging Model Considering TOU Price and SOC Curve». In: *IEEE Xplore* (2011).

[6] W. Sun, F. Neumann, and Harrison. «Robust scheduling of electric vehicle charging in LV distribution networks under uncertainty». In: *IEEE Transactions on Industry Applications* (2020).

[7] A. Di Giorgio, A. Di Maria, F. Liberati, V. Suraci, and F. Delli Priscoli. «Lagrangian Decomposition based Multi Agent Model Predictive Control for Electric Vehicles Charging integrating Real Time Pricing». In: *arXiv* (2016).

[8] S. Pirouzi. «Two alternative robust optimization models for flexible power management of electric vehicles in distribution networks». In: *Energy* (2017).

[9] T. D. Tran, N.-D. Nguyen, H. T. M. Chu, L. El Ghaoui, L. Ambrosino, and G. Calafiore. «A Robust Optimization Model for Cost-Efficient and Fast Electric Vehicle Charging with L2-norm Uncertainty». In: (2025).

[10] Y. Cao, L. Huang, Y. Li, K. Jermsittiparsert, H. Ahmadi-Nezamabad, and S. Nojavan. «Optimal scheduling of electric vehicles aggregator under market price uncertainty using robust optimization technique». In: *International Journal of Electrical Power & Energy Systems* (2020).

[11] A. Soroudi and A. Keane. «Robust optimization based EV charging». In: (2014).

[12] X. Xu, W. Hu, W. Liu, and Y. Du. «Robust energy management for an on-grid hybrid hydrogen refueling and battery swapping station based on renewable energy». In: (2022).

[13] B. Chen, X. C. Liu, R. Wei, J. Chen, and Z. Chen. «An agent-based modeling approach for public charging demand estimation and charging station location optimization at urban scale». In: (submited).

[14] J.-U. Yu, K.-S. Cho, S.-W. Park, and S.-Y. Son. «Digital Twin System Framework and Implementation for Grid-Integrated Electric Vehicles». In: ().

[15] L. Ambrosino, K. M. Nguyen, M. B. Vu, R. Zorgati, L. El Ghaoui, and G. C. Calafiore. «A Multi-Criterion Approach to Smart EV Charging with CO_2 Emissions and Cost Minimization». In: *Under Submission* (2025).

[16] G. C. Calafiore and L. Ambrosino. «A Cohort-Based Optimization Model for Electric Vehicles Charging». In: (2025). URL: https://doi.org/10.1016/j.ifacol.2025.08.130.

[17] EMBER. *European wholesale electricity price data.* https://ember-climate.org/data-catalogue/european-wholesale-electricity-price-data/. Accessed: 2025-08-11.

[18] Stephen Boyd and Lieven Vandenberghe. *Convex Optimization.* Cambridge University Press, 2004. URL: https://web.stanford.edu/~boyd/cvxbook/.

[19] Sébastien Bubeck. «Convex Optimization: Algorithms and Complexity». In: *Foundations and Trends in Machine Learning* 8.3-4 (2015), pp. 231–358.

[20] Aharon Ben-Tal and Arkadi Nemirovski. *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications.* SIAM, 2001.

[21] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. «AdWords and Generalized Online Matching». In: *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS).* 2007.

[22] Javad Lavaei and Steven H. Low. «Zero Duality Gap in Optimal Power Flow Problem». In: *IEEE Transactions on Power Systems* (2012).

[23] James B. Rawlings and David Q. Mayne. *Model Predictive Control: Theory and Design.* Nob Hill Publishing, 2009.

[24] Mung Chiang, Chih-Wei Tan, Daniel P. Palomar, Daniel O'Neill, and David Julian. «Geometric Programming for Communication Systems». In: *Foundations and Trends in Communications and Information Theory* 2 (2005).

[25] Michael Lustig, David Donoho, and John M. Pauly. «Sparse MRI: The Application of Compressed Sensing for Rapid MR Imaging». In: *Magnetic Resonance in Medicine* (2007).

[26] R. T. Marler and J. S. Arora. «Survey of multi-objective optimization methods for engineering». In: *Structural and Multidisciplinary Optimization* 26.6 (2004), pp. 369–395.

[27] Himanshu Jain and Kalyanmoy Deb. «Evolutionary multi-objective optimization: A review of the state of the art». In: *IEEE Transactions on Evolutionary Computation* 22.3 (2017), pp. 408–432.

[28] Gi-Hyoug Lee and Jiyong Kim. «Uncertainty analysis for electric vehicle charging demand and its impact on power distribution networks». In: *Energies* 12.3 (2019), p. 453.

[29] Joel Goh and Melvyn Sim. «Distributionally robust optimization: A review». In: *Operations Research* 58.4-part-1 (2010), pp. 902–917.

[30] John R. Birge and François Louveaux. *Introduction to Stochastic Programming*. Springer, 2011.

[31] Alexander Shapiro, Darinka Dentcheva, and Andrzej Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, 2009.

[32] Yuri Ermoliev and Roger Wets. «Stochastic Programming: Survey of Applications». In: *Annals of Operations Research* (1988).

[33] Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust Optimization*. Princeton University Press, 2009.

[34] Nicholas Metropolis and Stanislaw Ulam. «The Monte Carlo method». In: *Journal of the American Statistical Association* 44.247 (1949), pp. 335–341.

[35] GeeksforGeeks. *What is Monte Carlo Simulation?* Accessed: 2025-08-01. 2022. URL: https://www.geeksforgeeks.org/artificial-intelligence/what-is-monte-carlo-simulation/.

[36] Tito Homem-de-Mello and Güzin Bayraksan. «Monte Carlo Sampling-Based Methods for Stochastic Optimization». In: *Surveys in Operations Research and Management Science* (2014).

[37] Jiaxin Zhang. «Modern Monte Carlo Methods for Efficient Uncertainty Quantification and Propagation: A Survey». In: *arXiv preprint arXiv:2011.00680* (2020).

[38] Paul Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, 2003.

[39] Ramin Manouchehrynia, Salwani Abdullah, and Sarbani Singh Karam Singh. «Fatigue reliability assessment of an automobile coil spring under random strain loads using probabilistic technique». In: *Metals* (2020).

[40] Shabbir Ahmed and Alexander Shapiro. «Sampling-based methods for stochastic optimization: A comprehensive survey». In: *Mathematics of Operations Research* (2020).

[41] G. C. Calafiore, L. Ambrosino, K. M. Nguyen, R. Zorgati, D. Nguyen-Ngoc, and L. El Ghaoui. «Robust Power Scheduling for Smart Charging of Electric Vehicles». In: *Proceedings of the 23rd European Control Conference (ECC2025)* (2025).